

US011443232B1

(12) **United States Patent**
Rahman et al.

(10) **Patent No.:** **US 11,443,232 B1**
(45) **Date of Patent:** **Sep. 13, 2022**

(54) **ACTIVE LEARNING-BASED DATA LABELING SERVICE USING AN AUGMENTED MANIFEST**

USPC 707/609, 640, 737, 602, 736
See application file for complete search history.

(71) Applicant: **Amazon Technologies, Inc.**, Seattle, WA (US)

(56) **References Cited**

(72) Inventors: **Zahid Rahman**, Livermore, CA (US); **Wei Xiao**, Bellevue, WA (US); **Stefano Stefani**, Issaquah, WA (US); **Rahul Sharma**, San Jose, CA (US); **Siddharth Joshi**, Seattle, WA (US)

U.S. PATENT DOCUMENTS

7,792,353 B2 9/2010 Forman et al.
2011/0088011 A1* 4/2011 Ouali G06F 8/10
2012/0269436 A1 10/2012 Mensink et al. 717/105
(Continued)

(73) Assignee: **Amazon Technologies, Inc.**, Seattle, WA (US)

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 418 days.

Non Final Office Action, U.S. Appl. No. 16/370,706, dated Sep. 1, 2020, 8 pages.

(Continued)

(21) Appl. No.: **16/370,733**

Primary Examiner — Md I Uddin

(22) Filed: **Mar. 29, 2019**

(74) *Attorney, Agent, or Firm* — Nicholson De Vos Webster & Elliott LLP

Related U.S. Application Data

(60) Provisional application No. 62/770,882, filed on Nov. 23, 2018.

(51) **Int. Cl.**
G06F 7/00 (2006.01)
G06N 20/00 (2019.01)
G06F 16/335 (2019.01)
G06F 16/383 (2019.01)
H04L 67/10 (2022.01)
G06F 40/40 (2020.01)

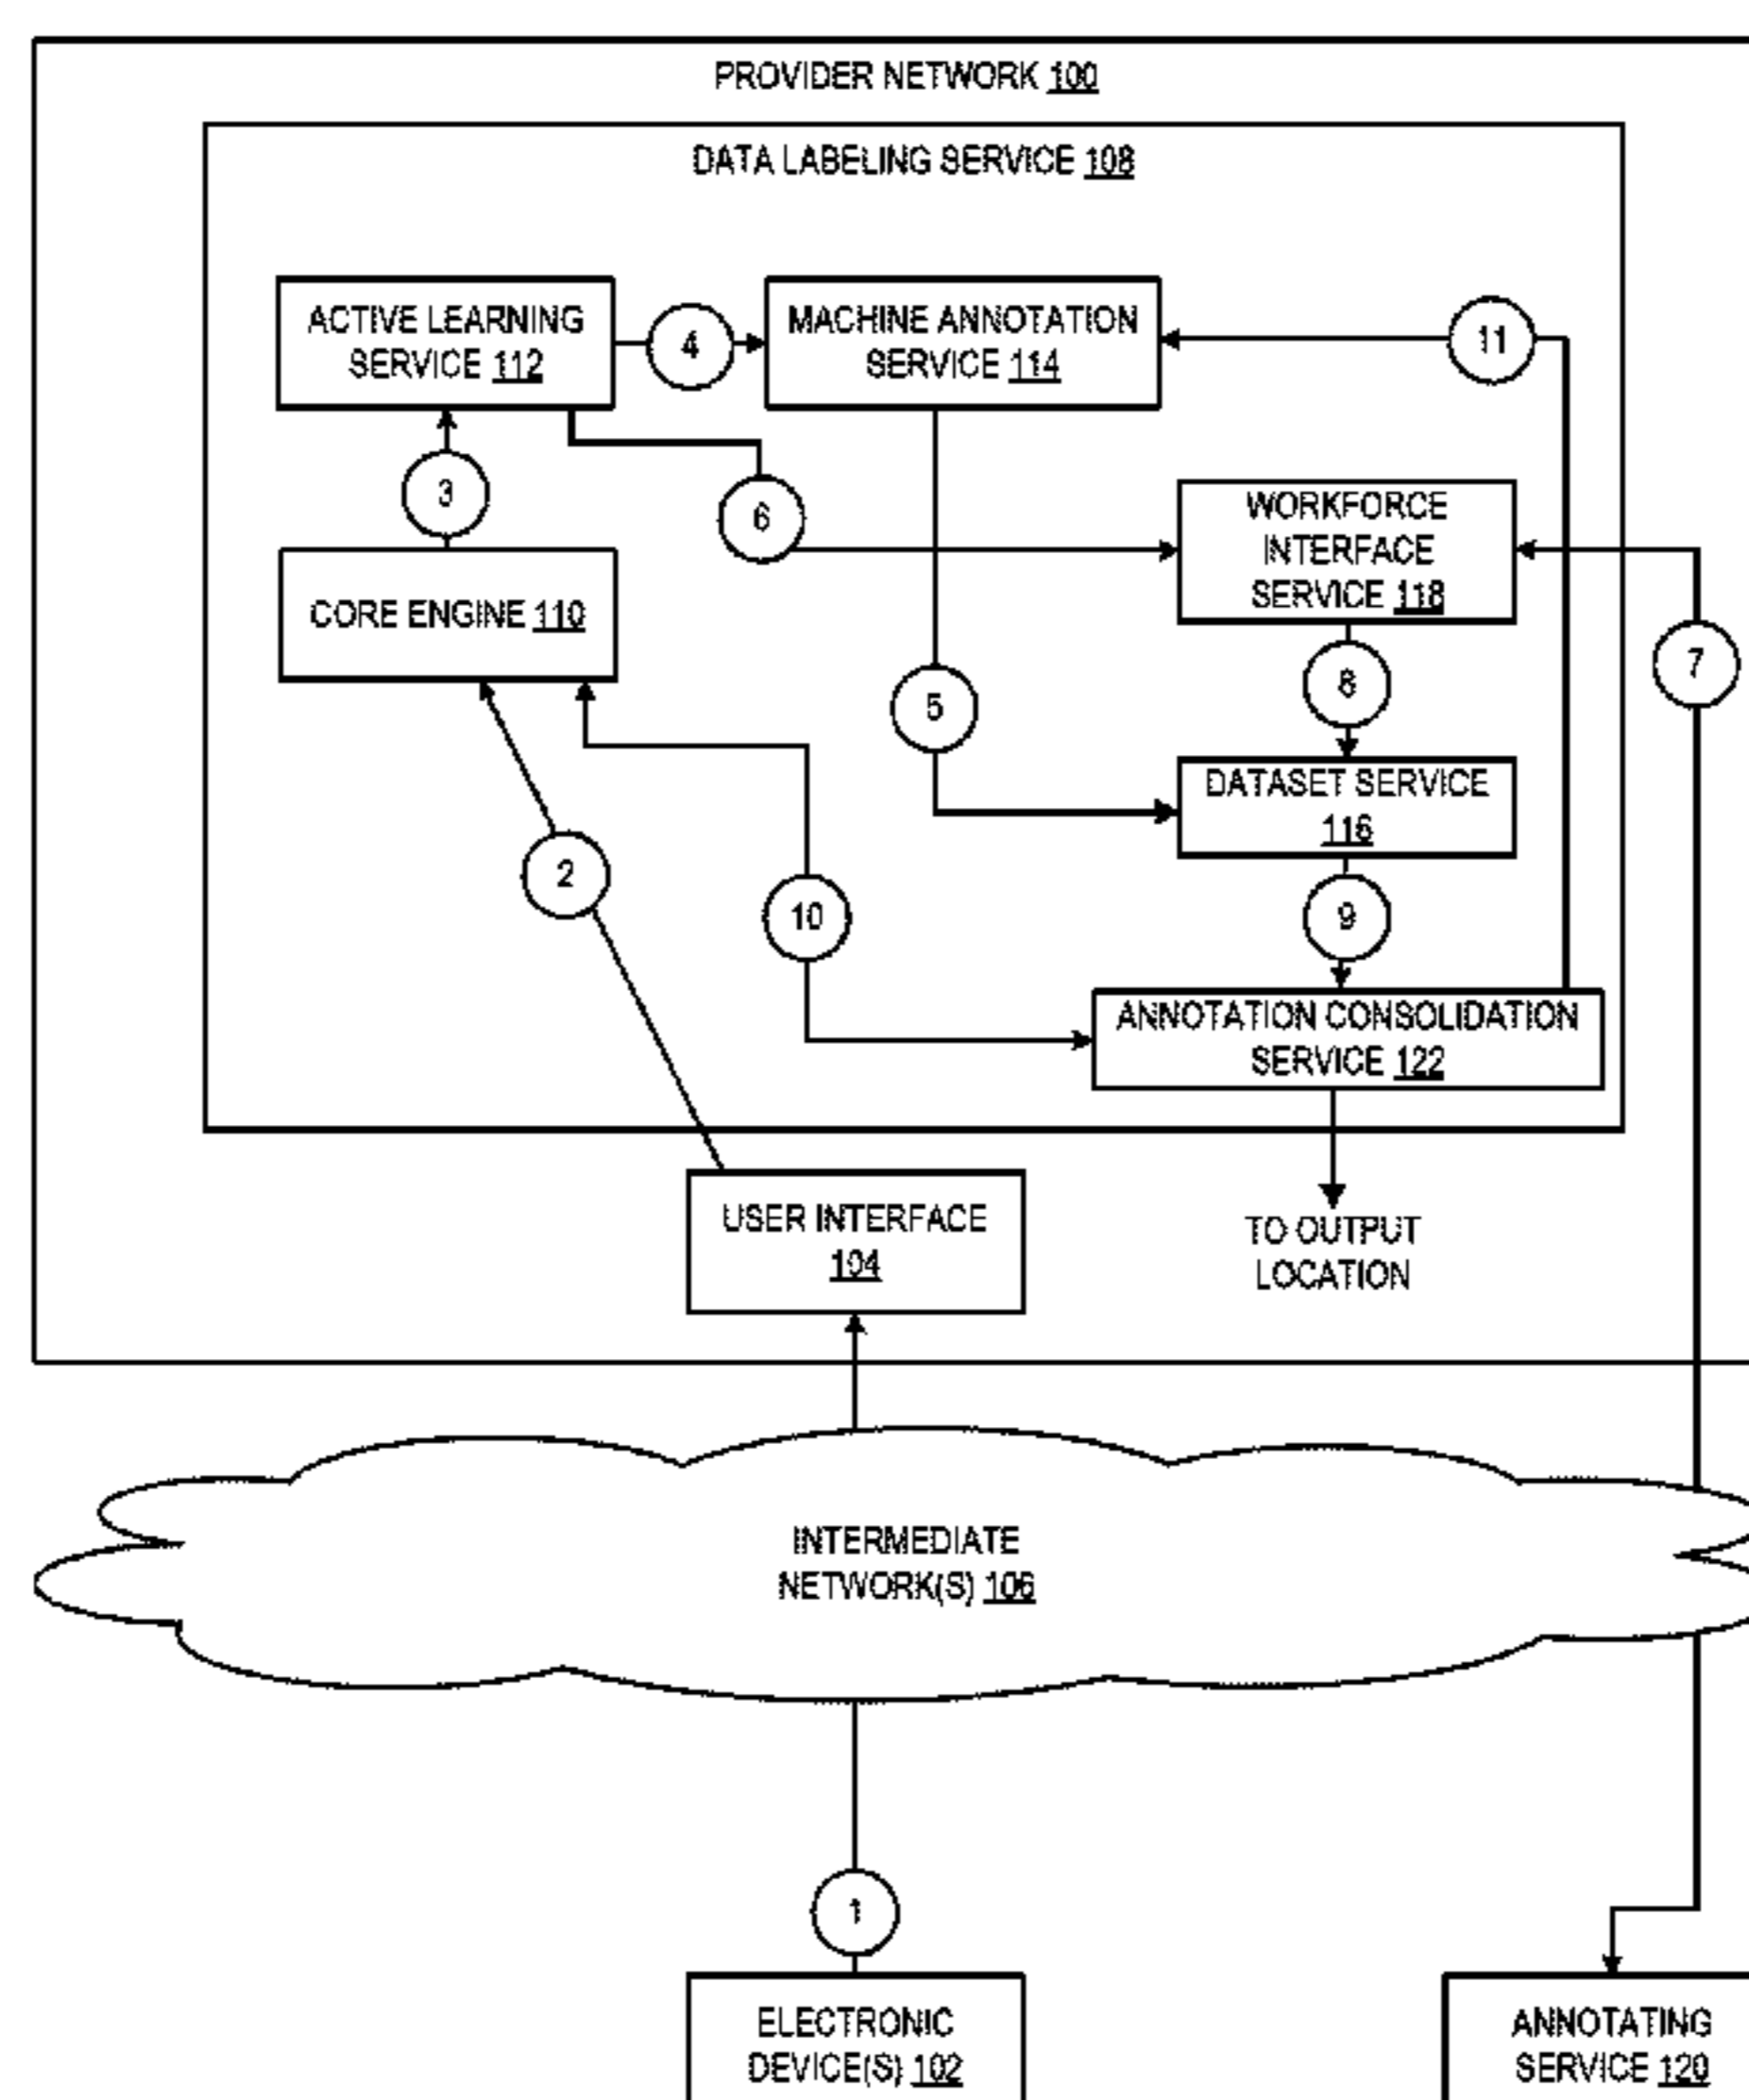
(57) **ABSTRACT**

Techniques for active learning-based data labeling are described. An active learning-based data labeling service enables a user to build and manage large, high accuracy datasets for use in various machine learning systems. Machine learning may be used to automate annotation and management of the datasets, increasing efficiency of labeling tasks and reducing the time required to perform labeling. Embodiments utilize active learning techniques to reduce the amount of a dataset that requires manual labeling. As subsets of the dataset are labeled, this label data is used to train a model which can then identify additional objects in the dataset without manual intervention. The label data can be added to an augmented manifest, the augmented manifest can be used to filter the dataset to perform further labeling jobs on the same or different subsets of the dataset.

(52) **U.S. Cl.**
CPC **G06N 20/00** (2019.01); **G06F 16/335** (2019.01); **G06F 16/383** (2019.01); **G06F 40/40** (2020.01); **H04L 67/10** (2013.01)

(58) **Field of Classification Search**
CPC G06F 16/152; G06F 16/137; G06F 16/184; G06F 16/335; G06F 16/383; G06F 40/40

20 Claims, 10 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2017/0039038 A1* 2/2017 Huber G06F 11/3072
2017/0315979 A1 11/2017 Boucher et al.
2018/0052891 A1 2/2018 Shuster et al.
2019/0028431 A1 1/2019 Keller et al.
2019/0294881 A1 9/2019 Polak et al.
2020/0226431 A1 7/2020 Zeiler et al.
2020/0401851 A1 12/2020 Mau et al.

OTHER PUBLICATIONS

Requirement for Restriction/Election, U.S. Appl. No. 16/370,723,
dated Aug. 10, 2020, 7 pages.

Non-Final Rejection U.S. Appl. No. 16/370,723, dated Feb. 5, 2021,
9 pages.

Notice of Allowance, U.S. Appl. No. 16/370,706, dated Mar. 2,
2021, 10 pages.

* cited by examiner

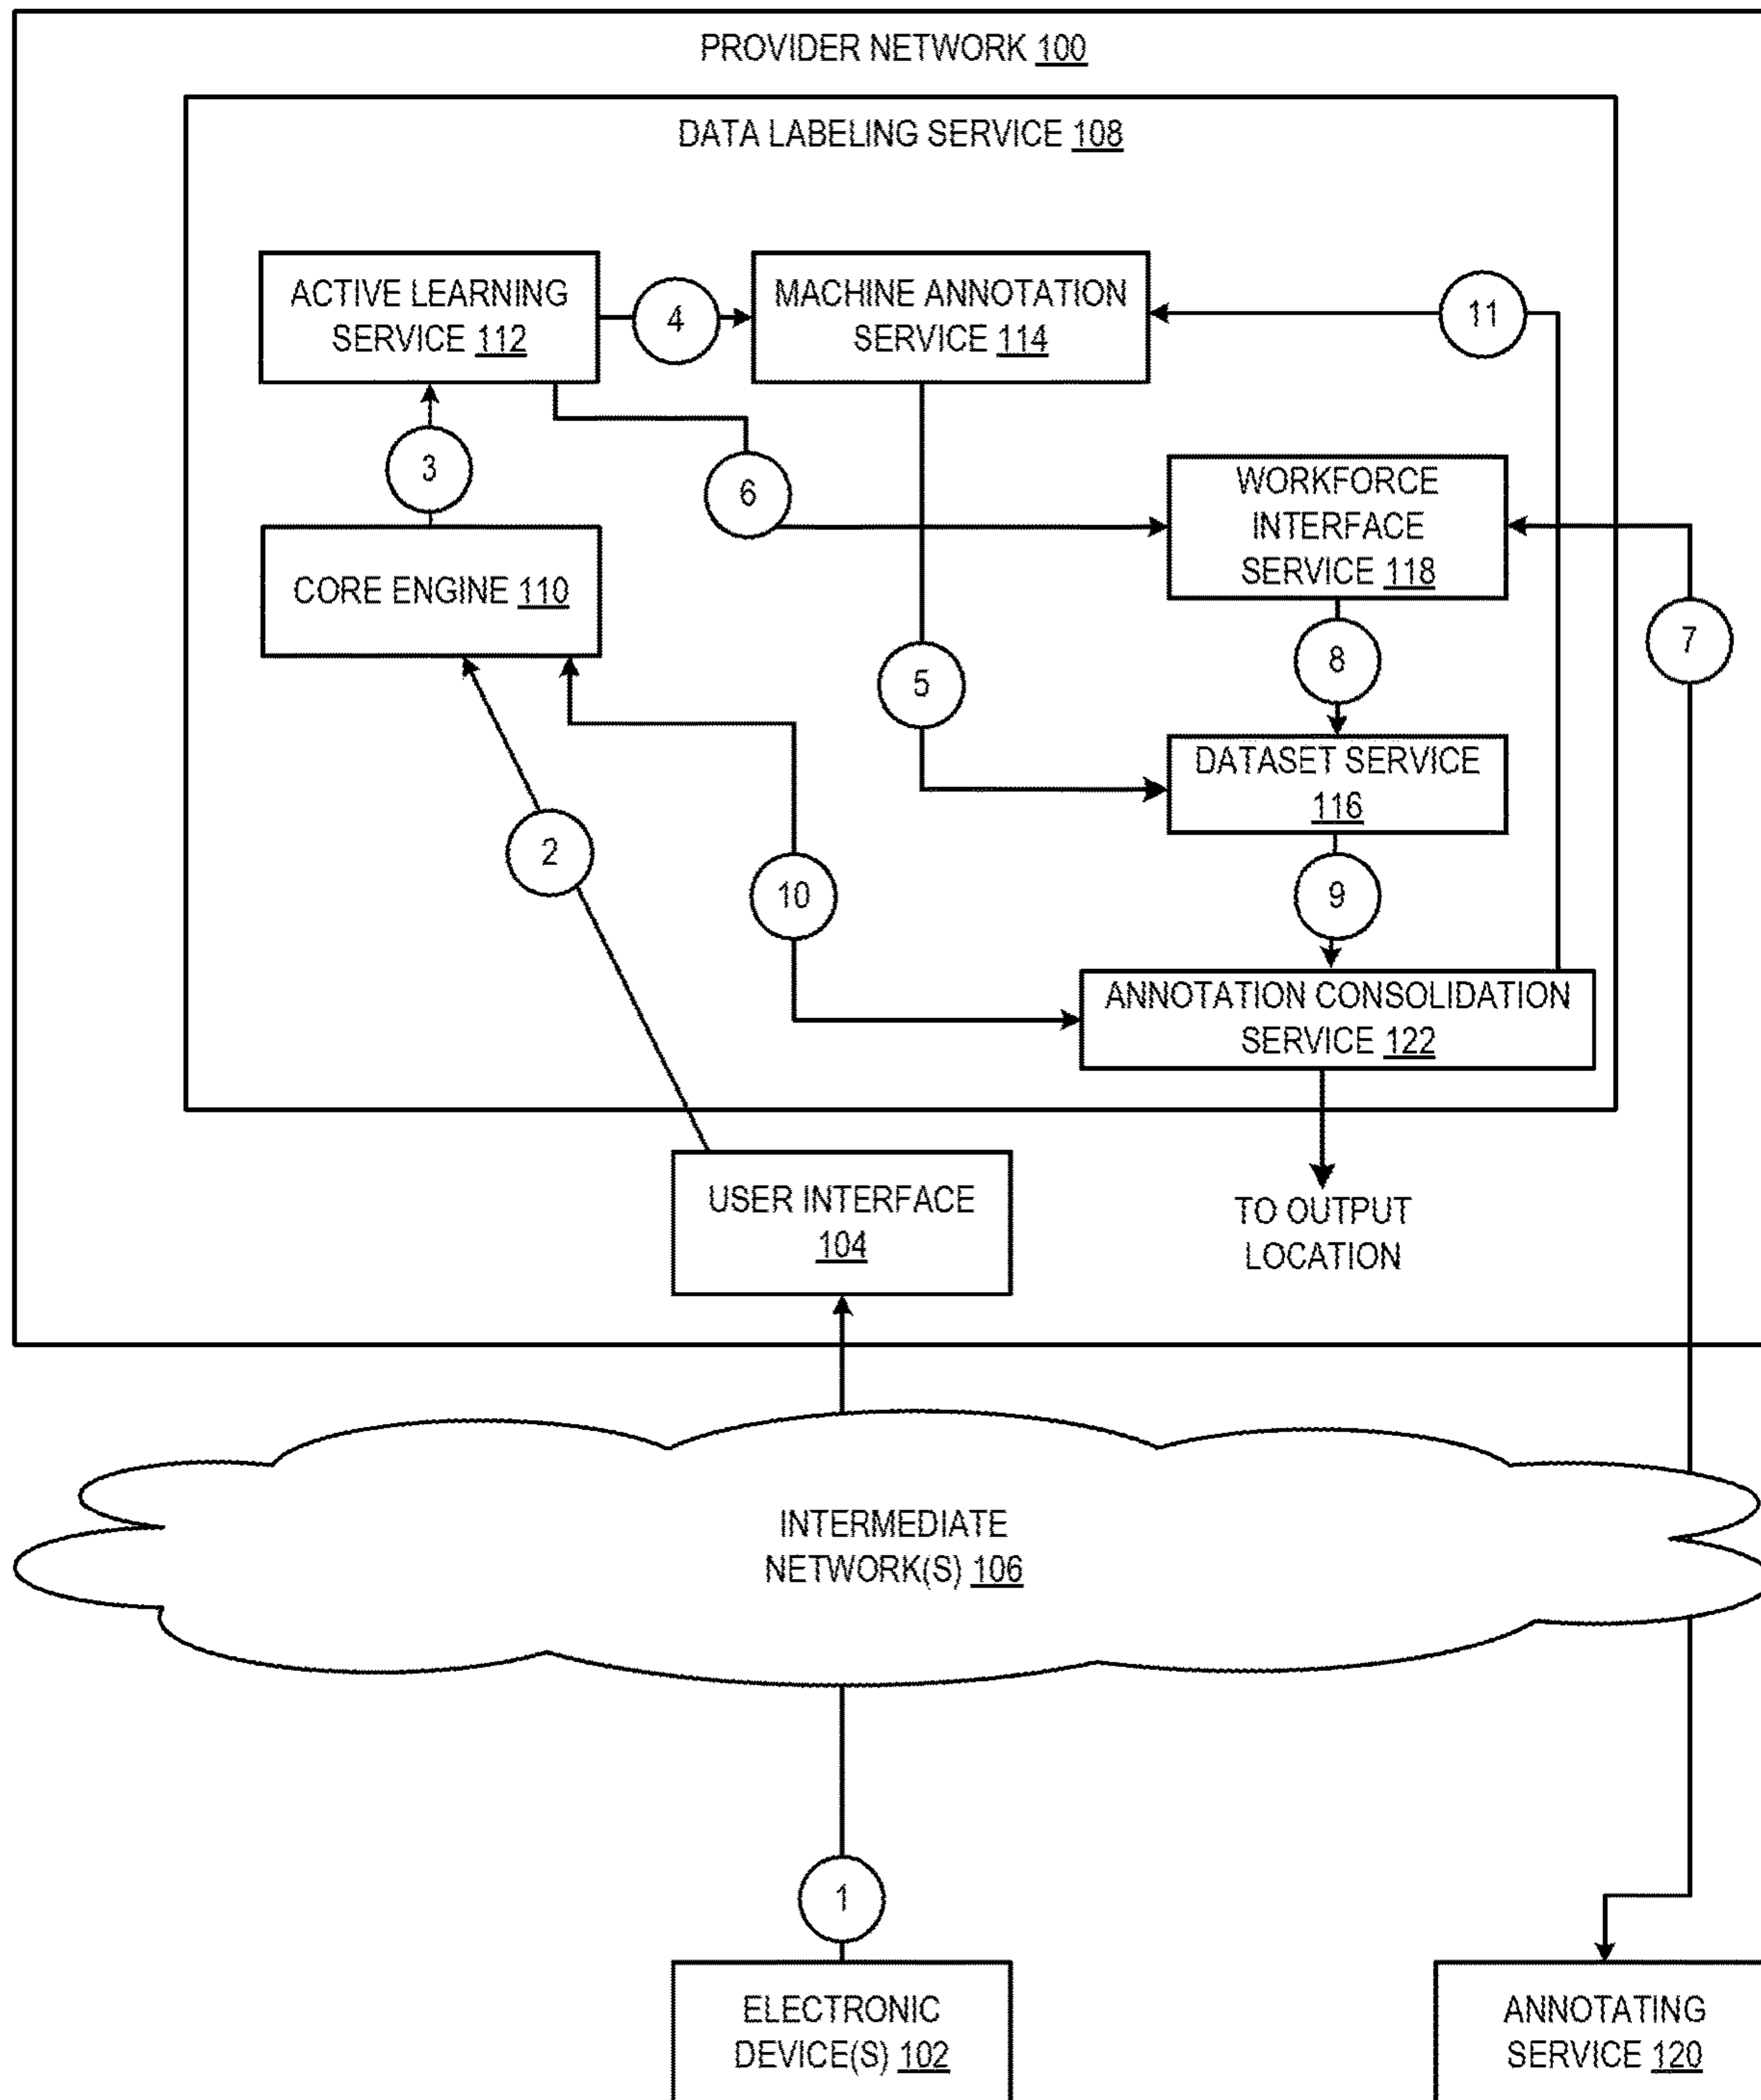


FIG. 1

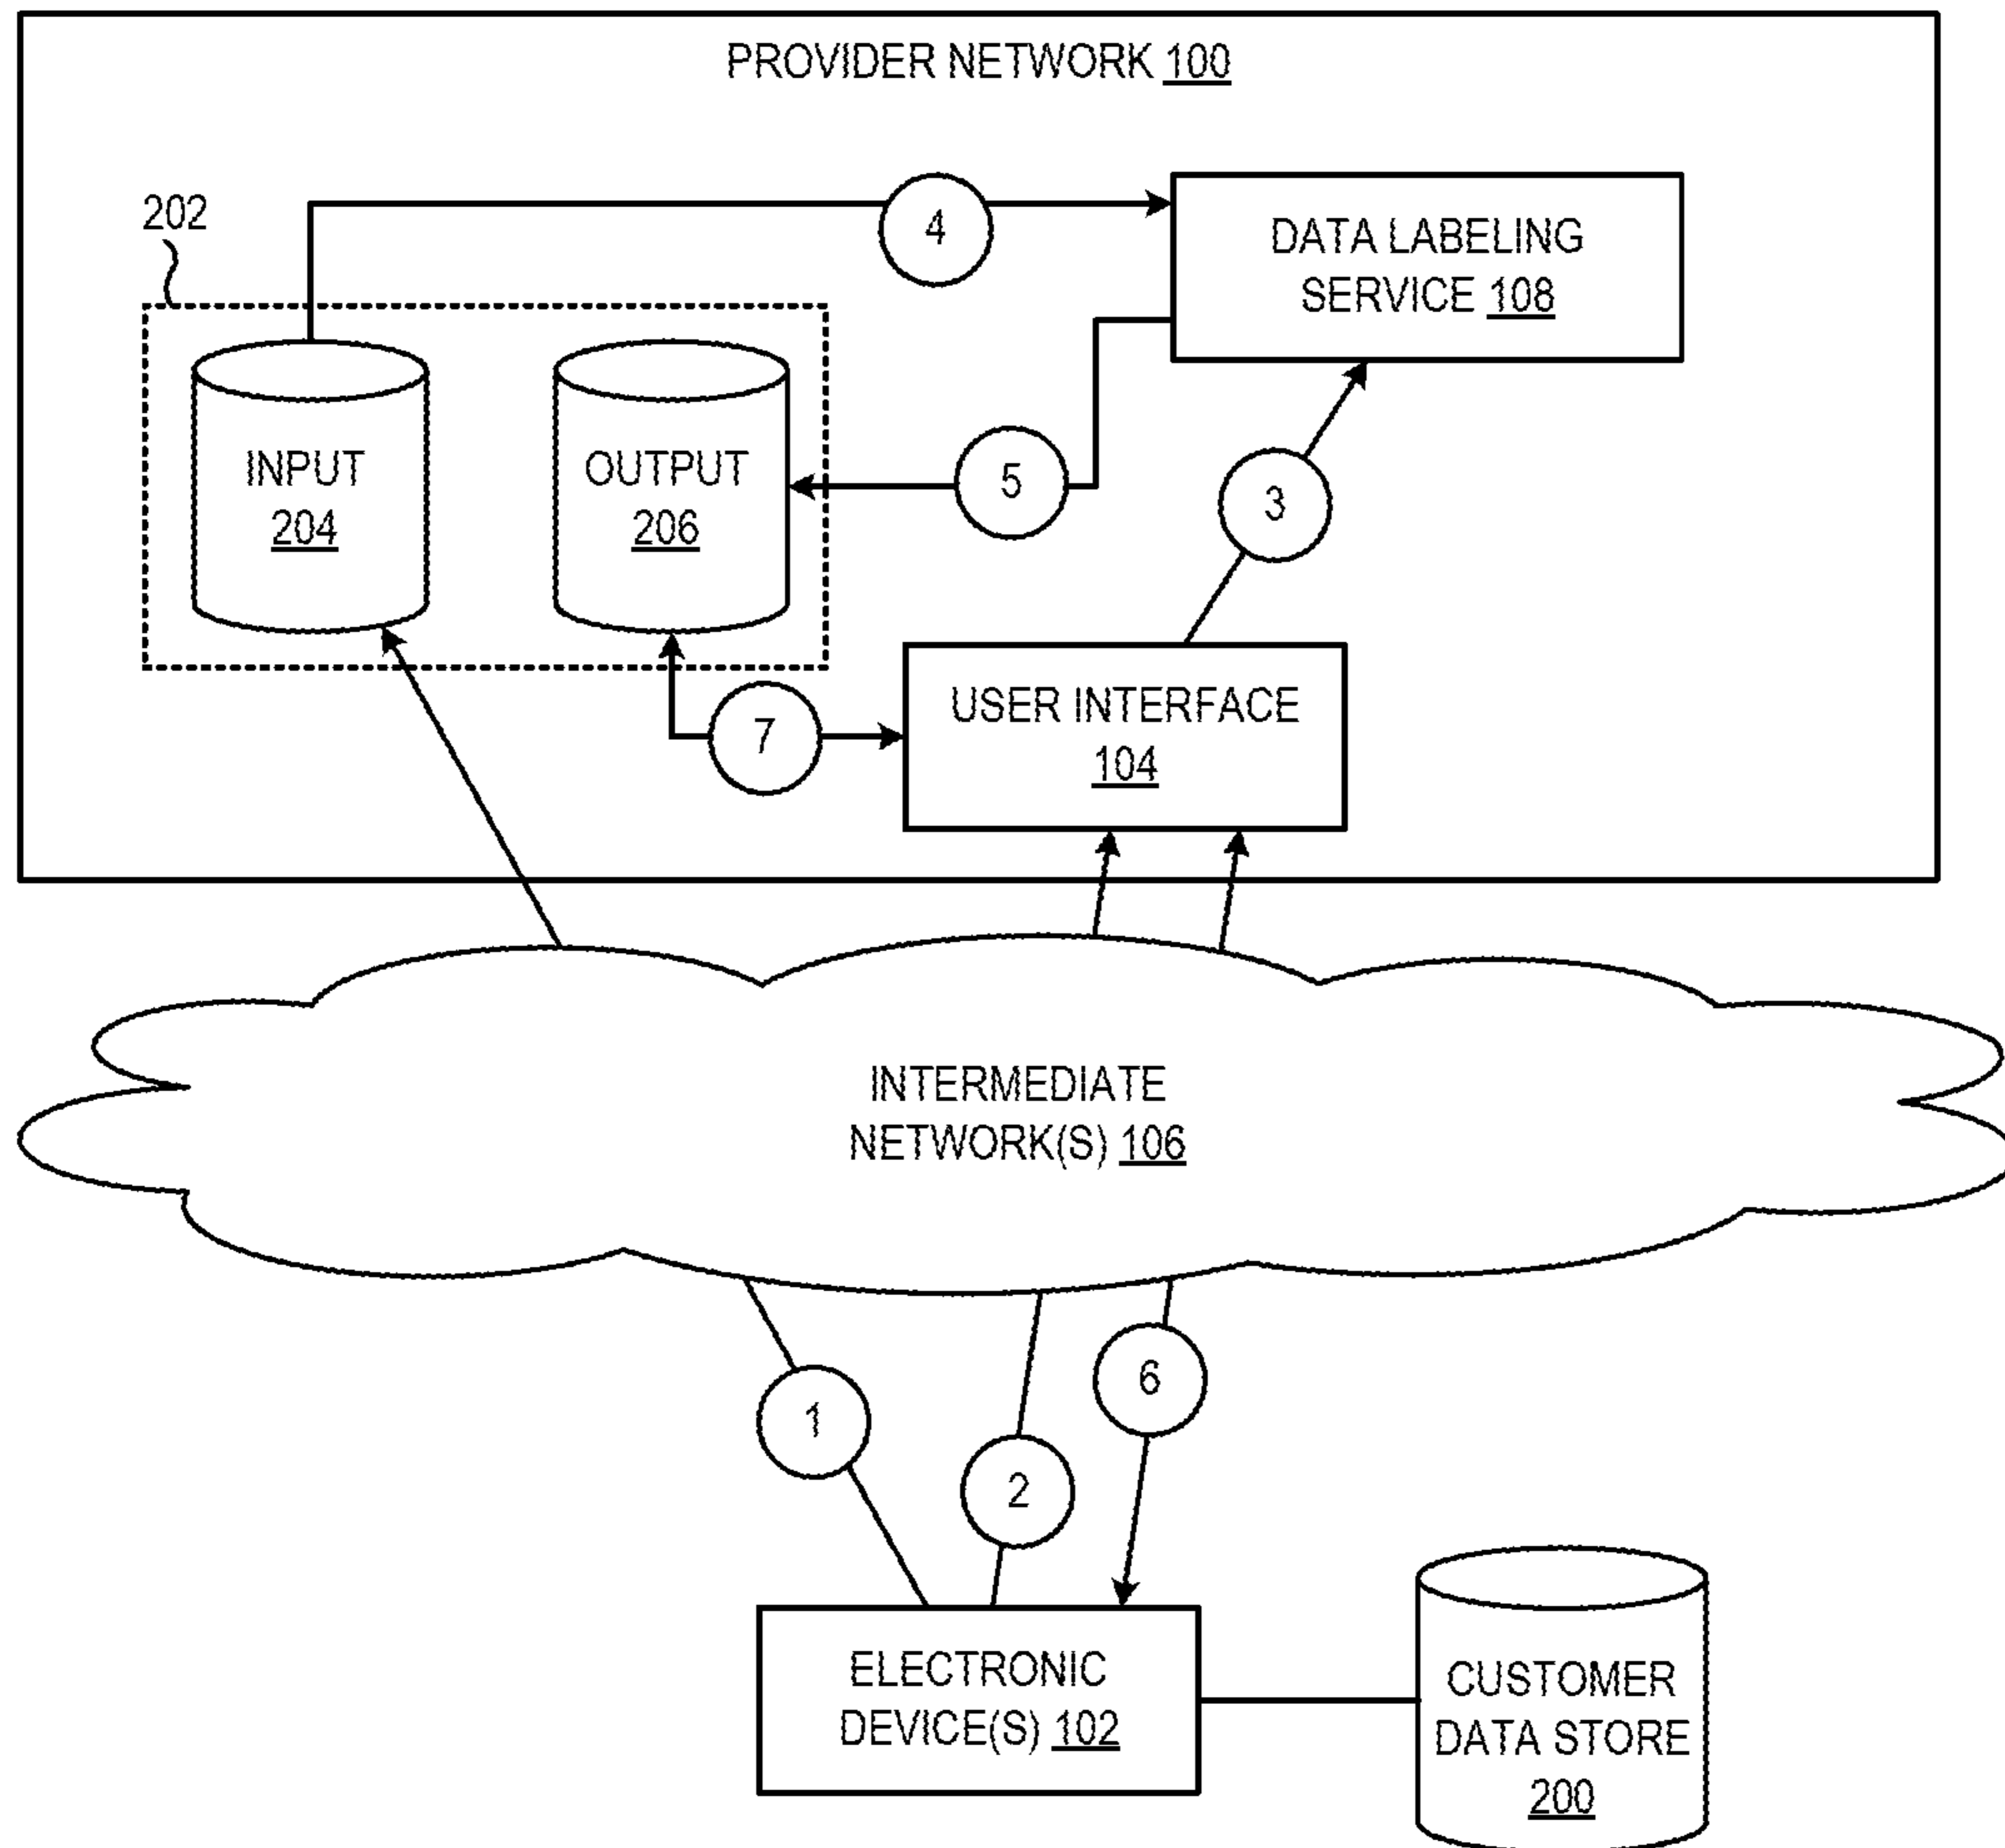


FIG. 2

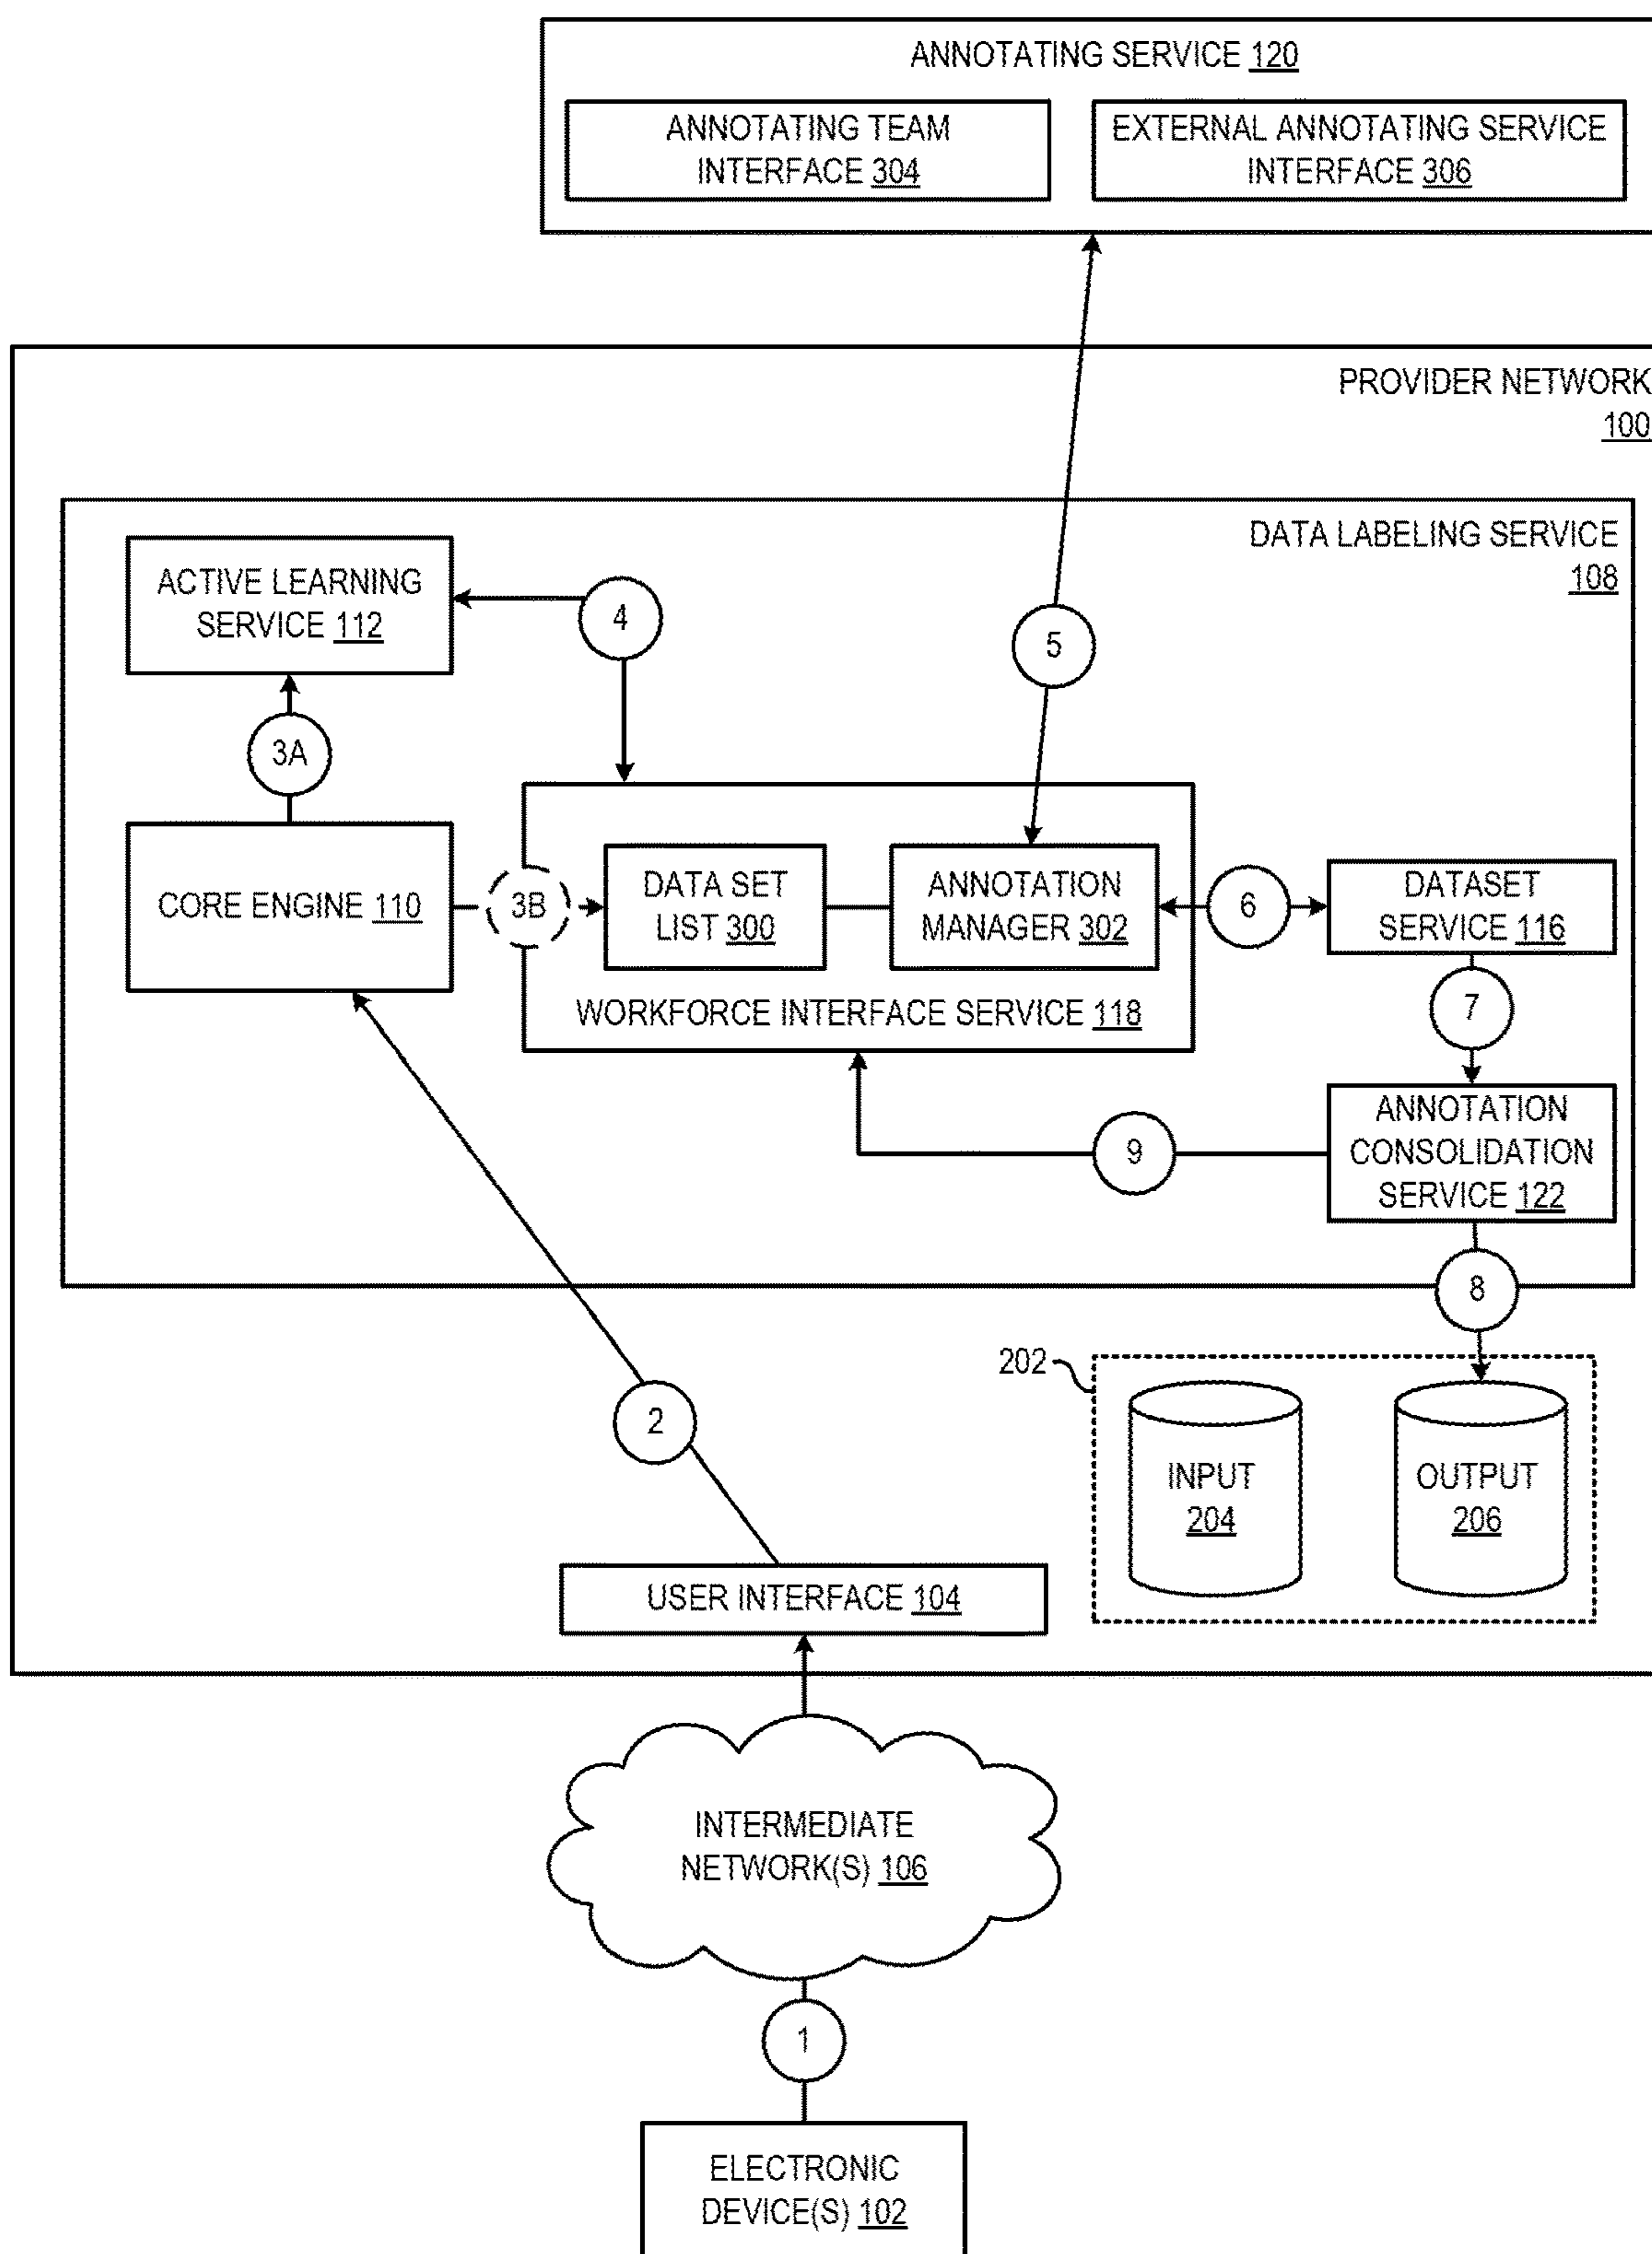


FIG. 3

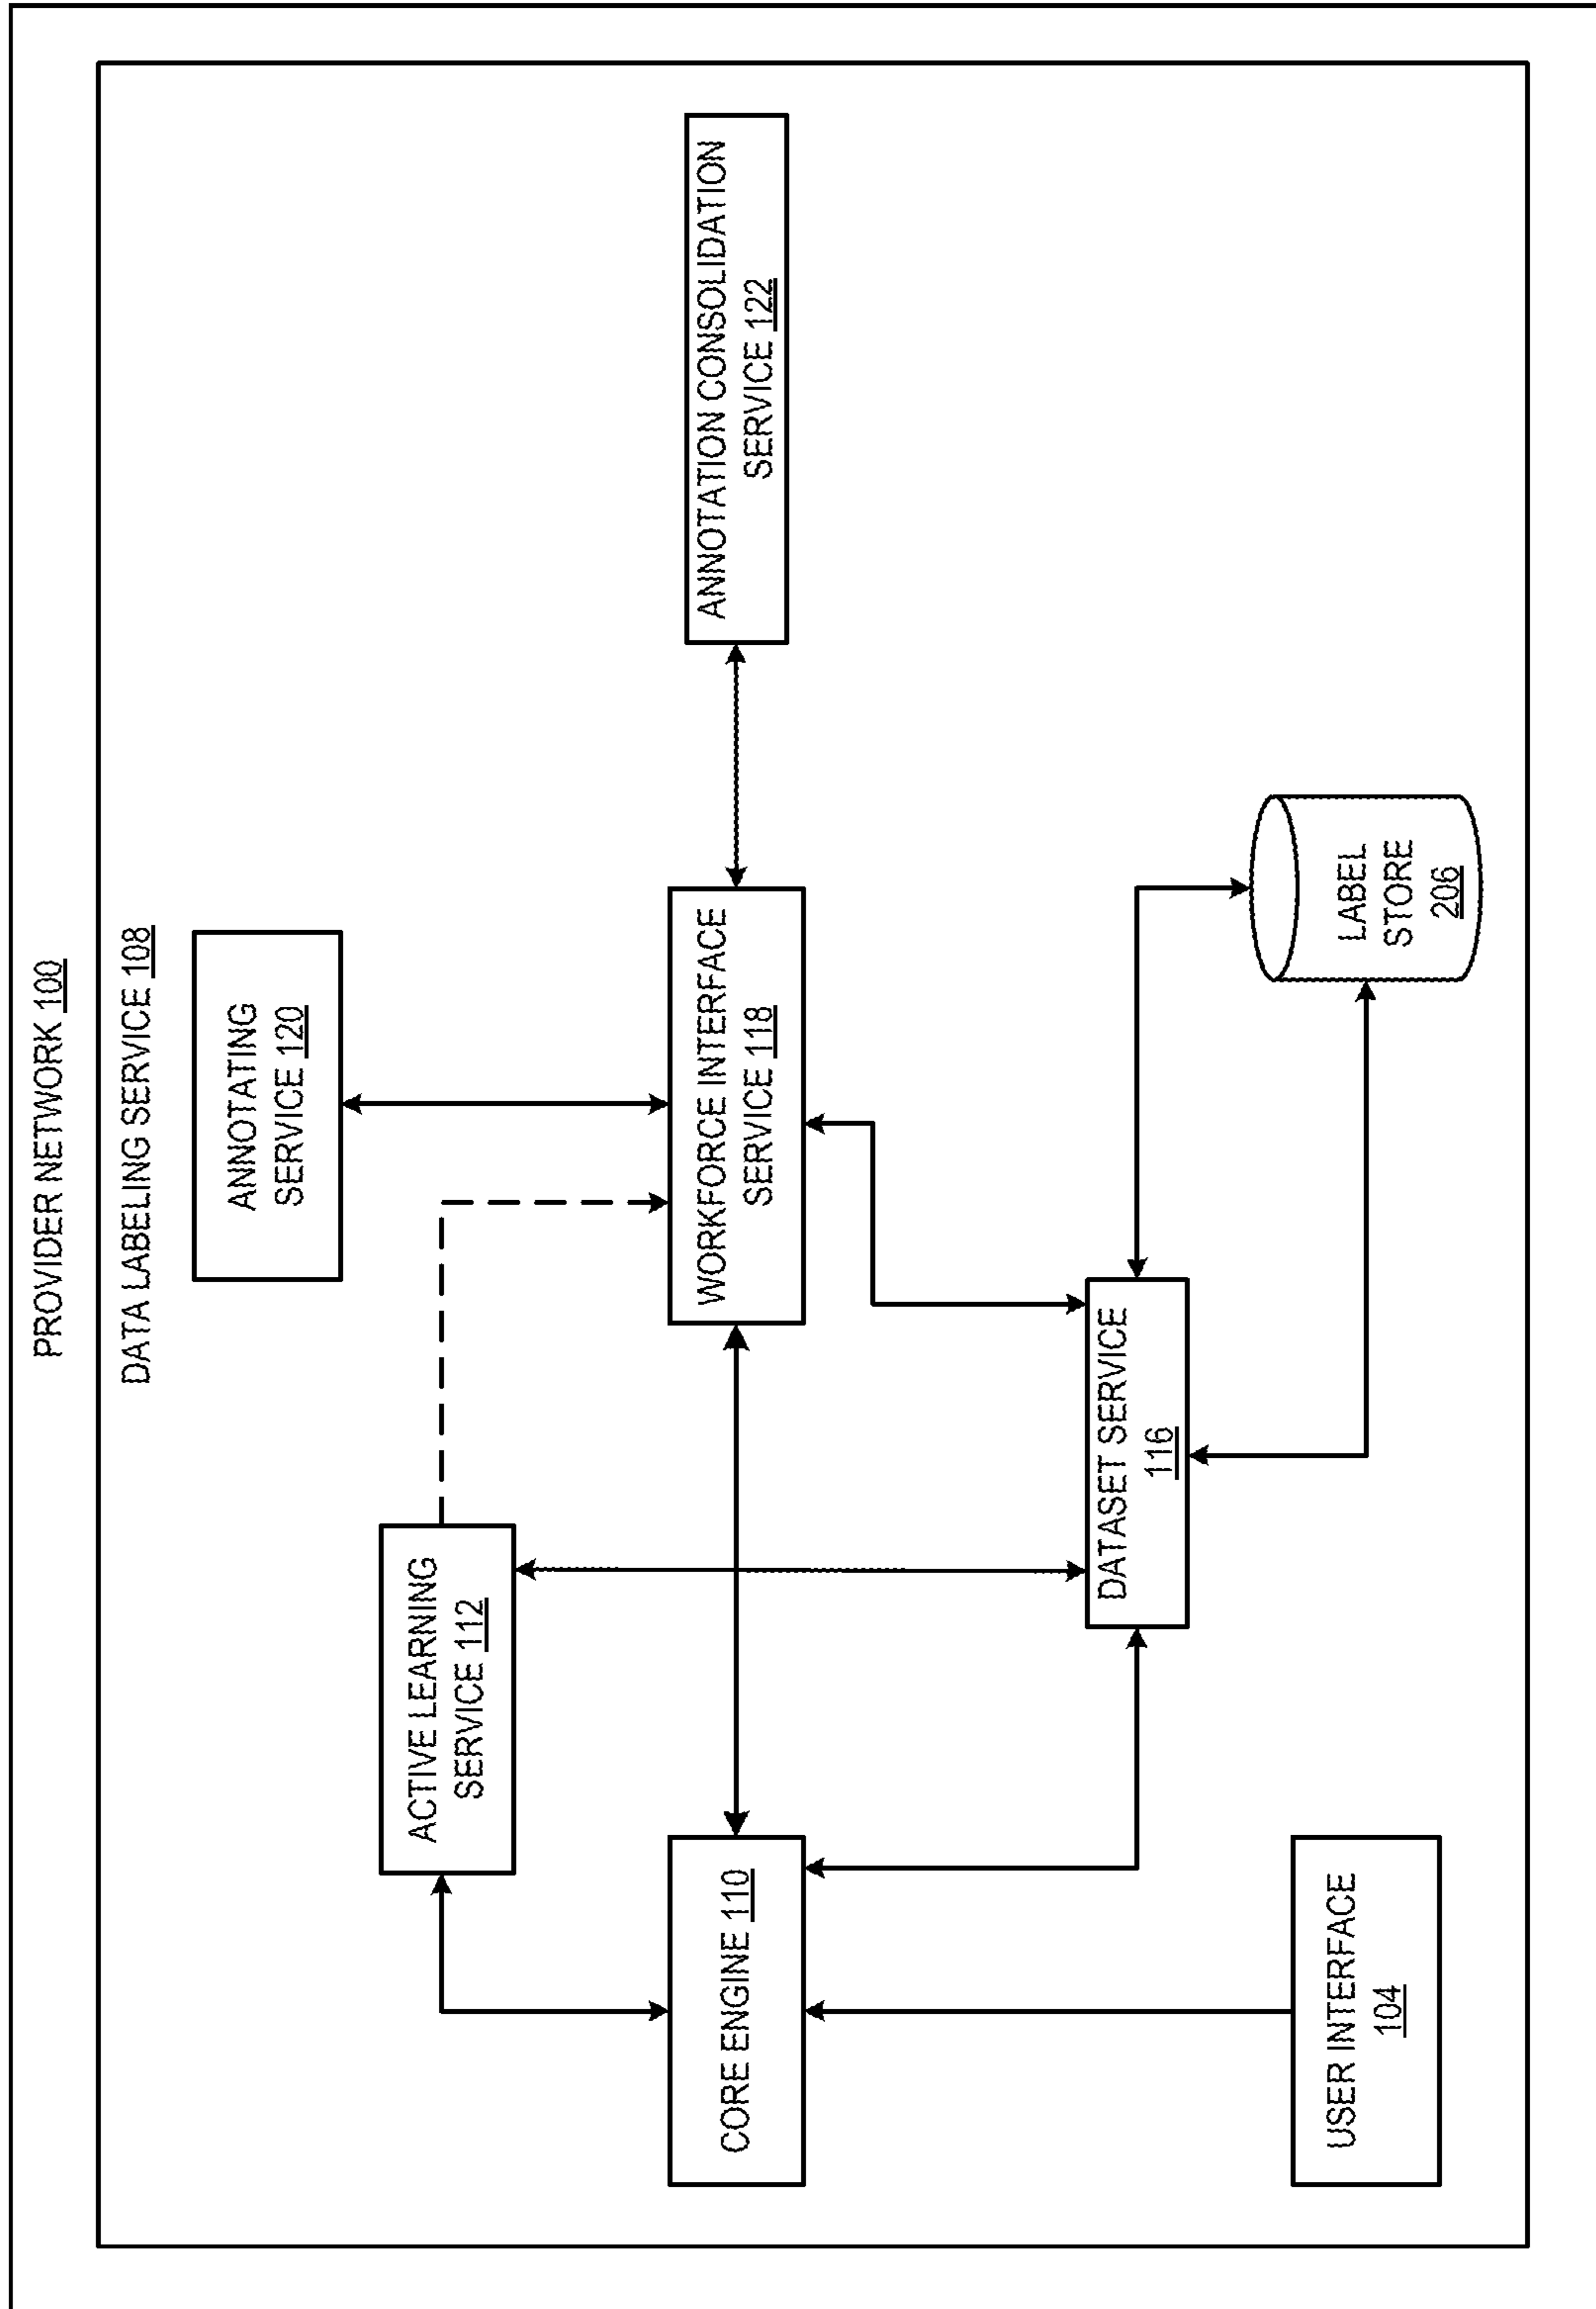


FIG. 4

500

```

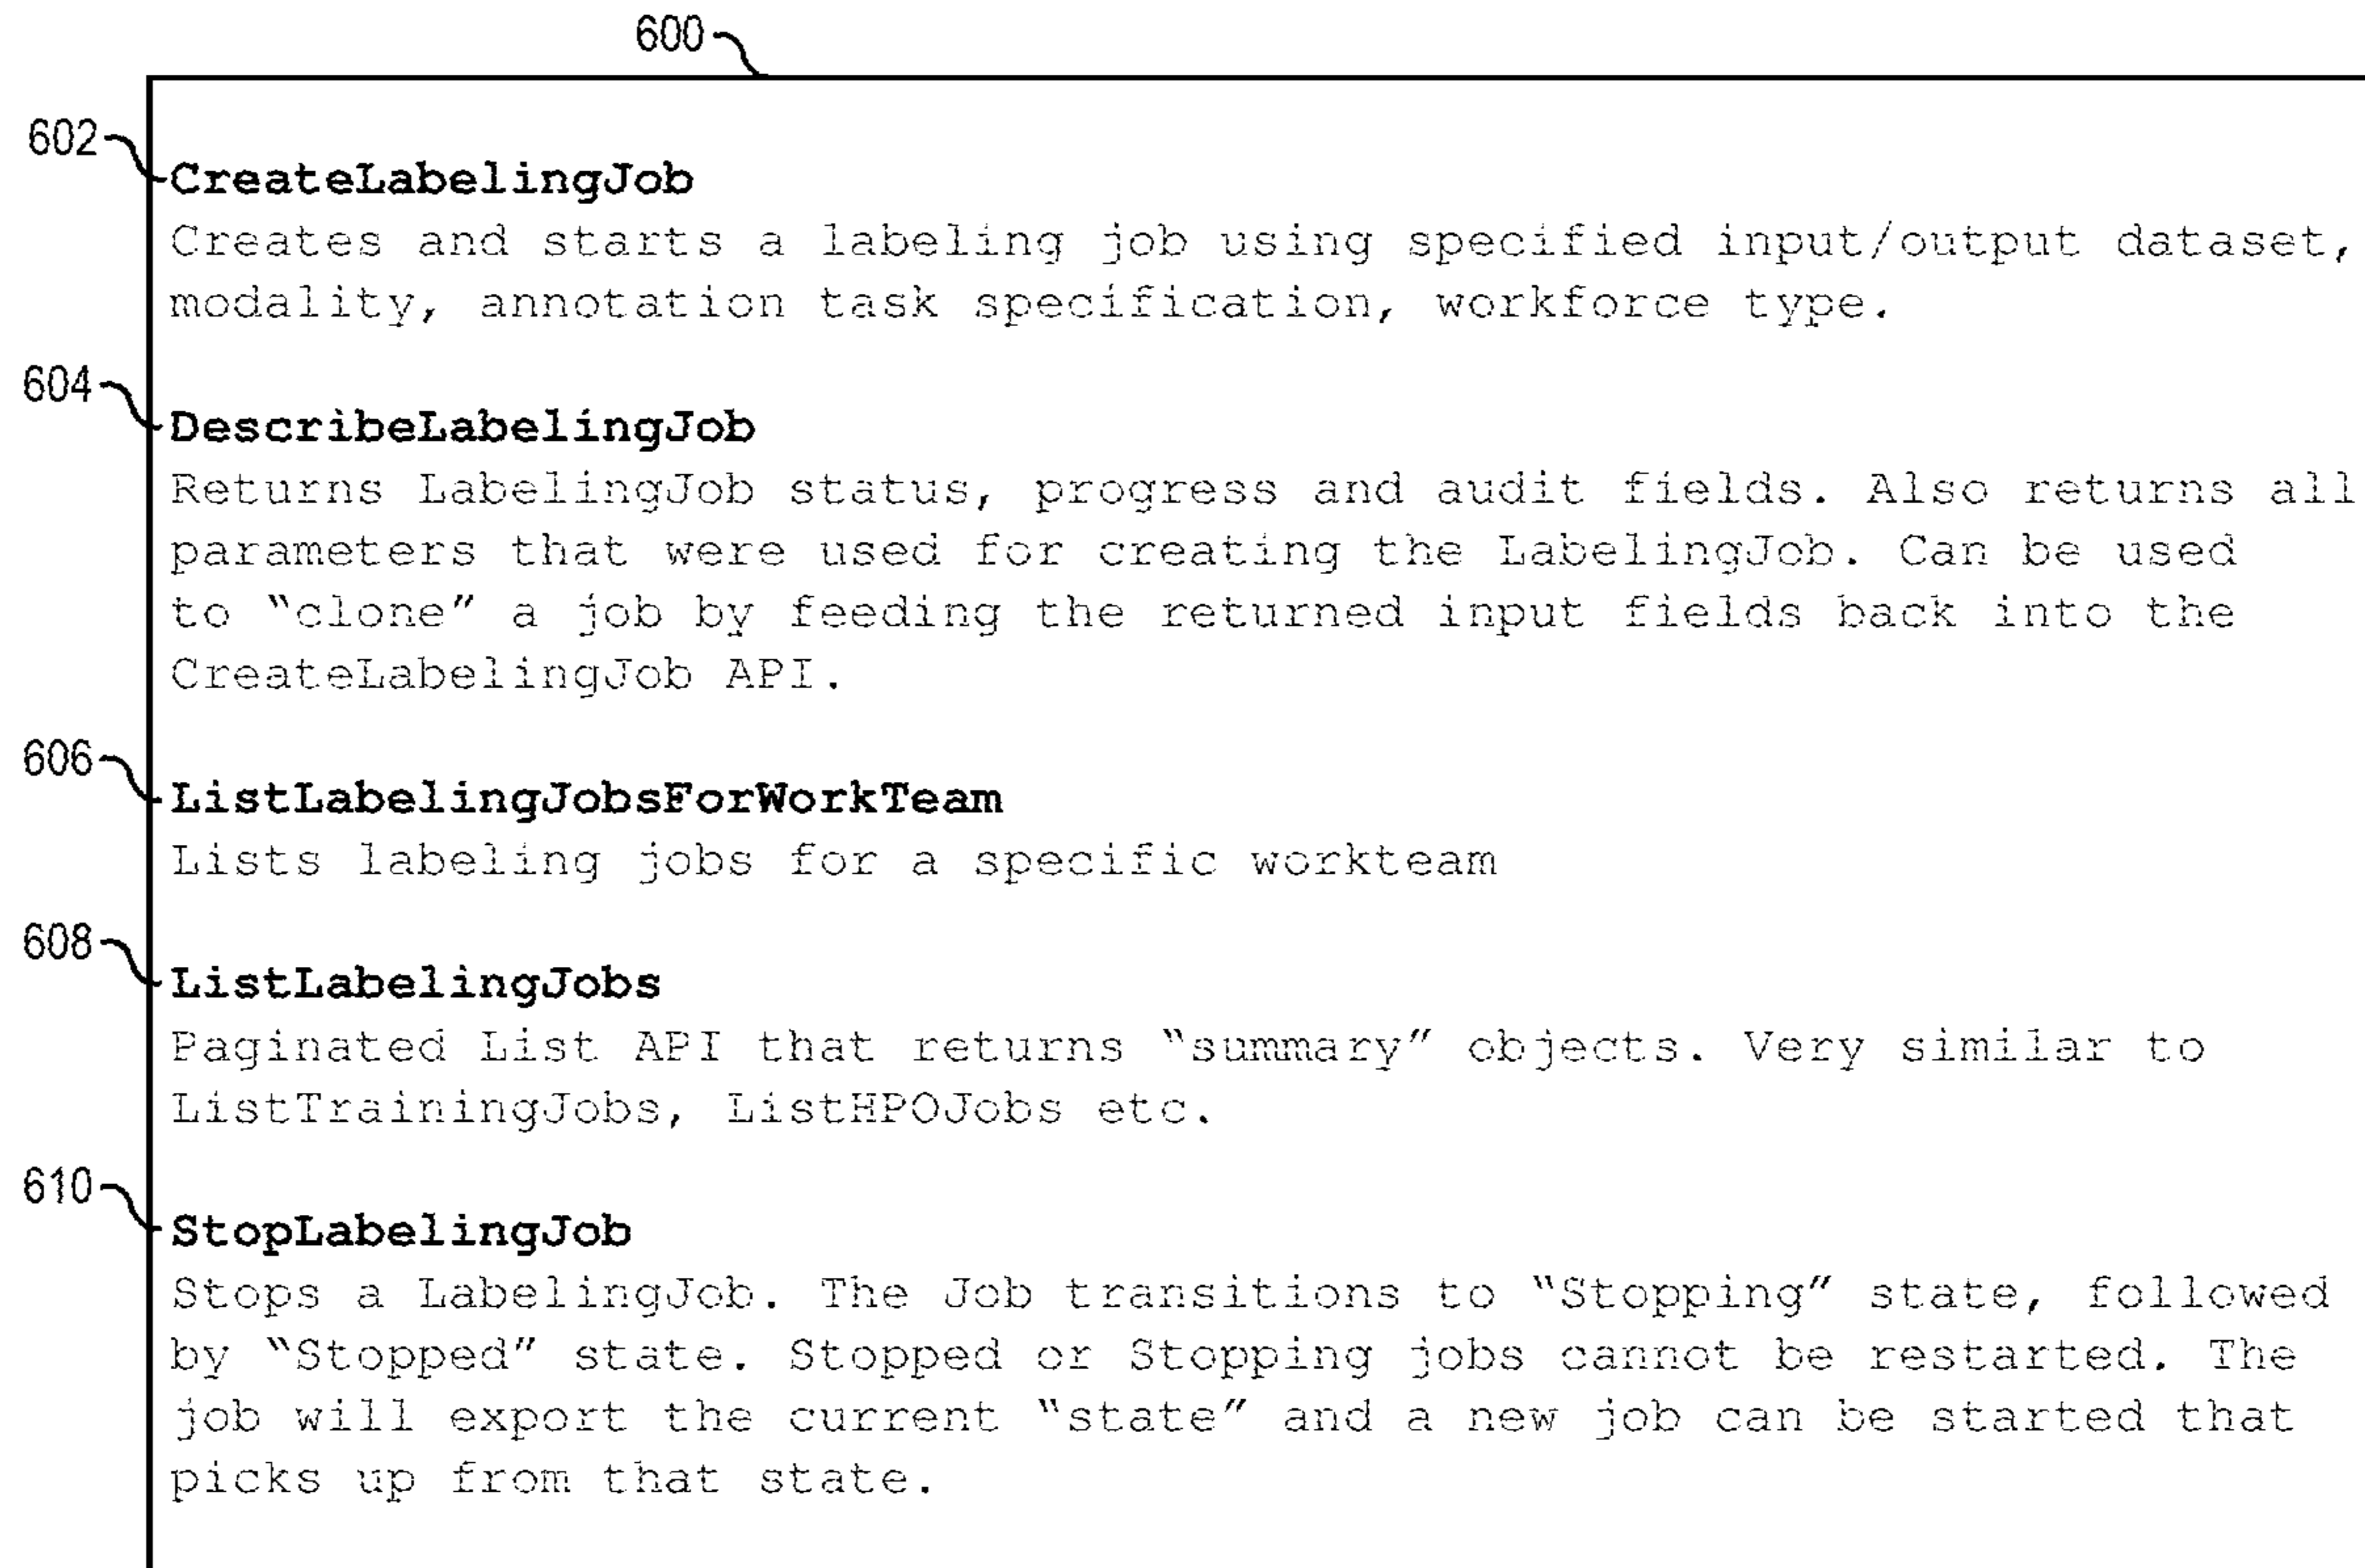
{"source-ref": "s3://cnidus-ml-iad/datasets/streetscenes/SSDB00000.JPG"}

{"source-ref": "s3://cnidus-ml-iad/datasets/streetscenes/SSDB00001.JPG", "streetscenes-road-objects1-
metadata": {"confidence": 0.95, "job-name": "labeling-job/streetscenes-road-objects1", "class-name": "vehicles", "human-
annotated": "yes", "creation-date": "2018-12-12T01:30:14.449763", "type": "groundtruth/image-classification"}}

{"source-ref": "s3://cnidus-ml-iad/datasets/streetscenes/SSDB00002.JPG", "streetscenes-road-objects1-
metadata": {"confidence": 0.95, "job-name": "labeling-job/streetscenes-road-objects1", "class-name": "vehicles", "human-
annotated": "yes", "creation-date": "2018-12-12T01:30:14.449763", "type": "groundtruth/image-classification"},
"streetscenes-road-objects2": [{"class_id": 0, "width": 38, "top": 490, "height": 24, "left": 351},
{"class_id": 0, "width": 85, "top": 505, "height": 53, "left": 450}, {"class_id": 0, "width": 65, "top": 489, "height": 43, "left": 592},
{"class_id": 0, "width": 59, "top": 480, "height": 43, "left": 524}, {"class_id": 0, "width": 354, "top": 471, "height": 150, "left": 567}],
"image_size": [{"width": 1280, "depth": 3, "height": 960}], "streetscenes-road-objects2-metadata": {"job-name": "labeling-job/streetscenes-road-
objects2", "class-map": {"0": "car"}, "human-annotated": "yes", "objects": [{"confidence": 0.09}, {"confidence": 0.09},
{"confidence": 0.09}], "creation-date": "2018-12-12T18:40:15.710919", "type": "groundtruth/object-detection"}

{"source-ref": "s3://cnidus-ml-iad/datasets/streetscenes/SSDB00003.JPG", "streetscenes-road-objects1-
metadata": {"confidence": 0.95, "job-name": "labeling-job/streetscenes-road-objects1", "class-name": "vehicles", "human-
annotated": "yes", "creation-date": "2018-12-12T01:21:57.370330", "type": "groundtruth/image-classification"}}
    
```

FIG. 5

**FIG. 6**

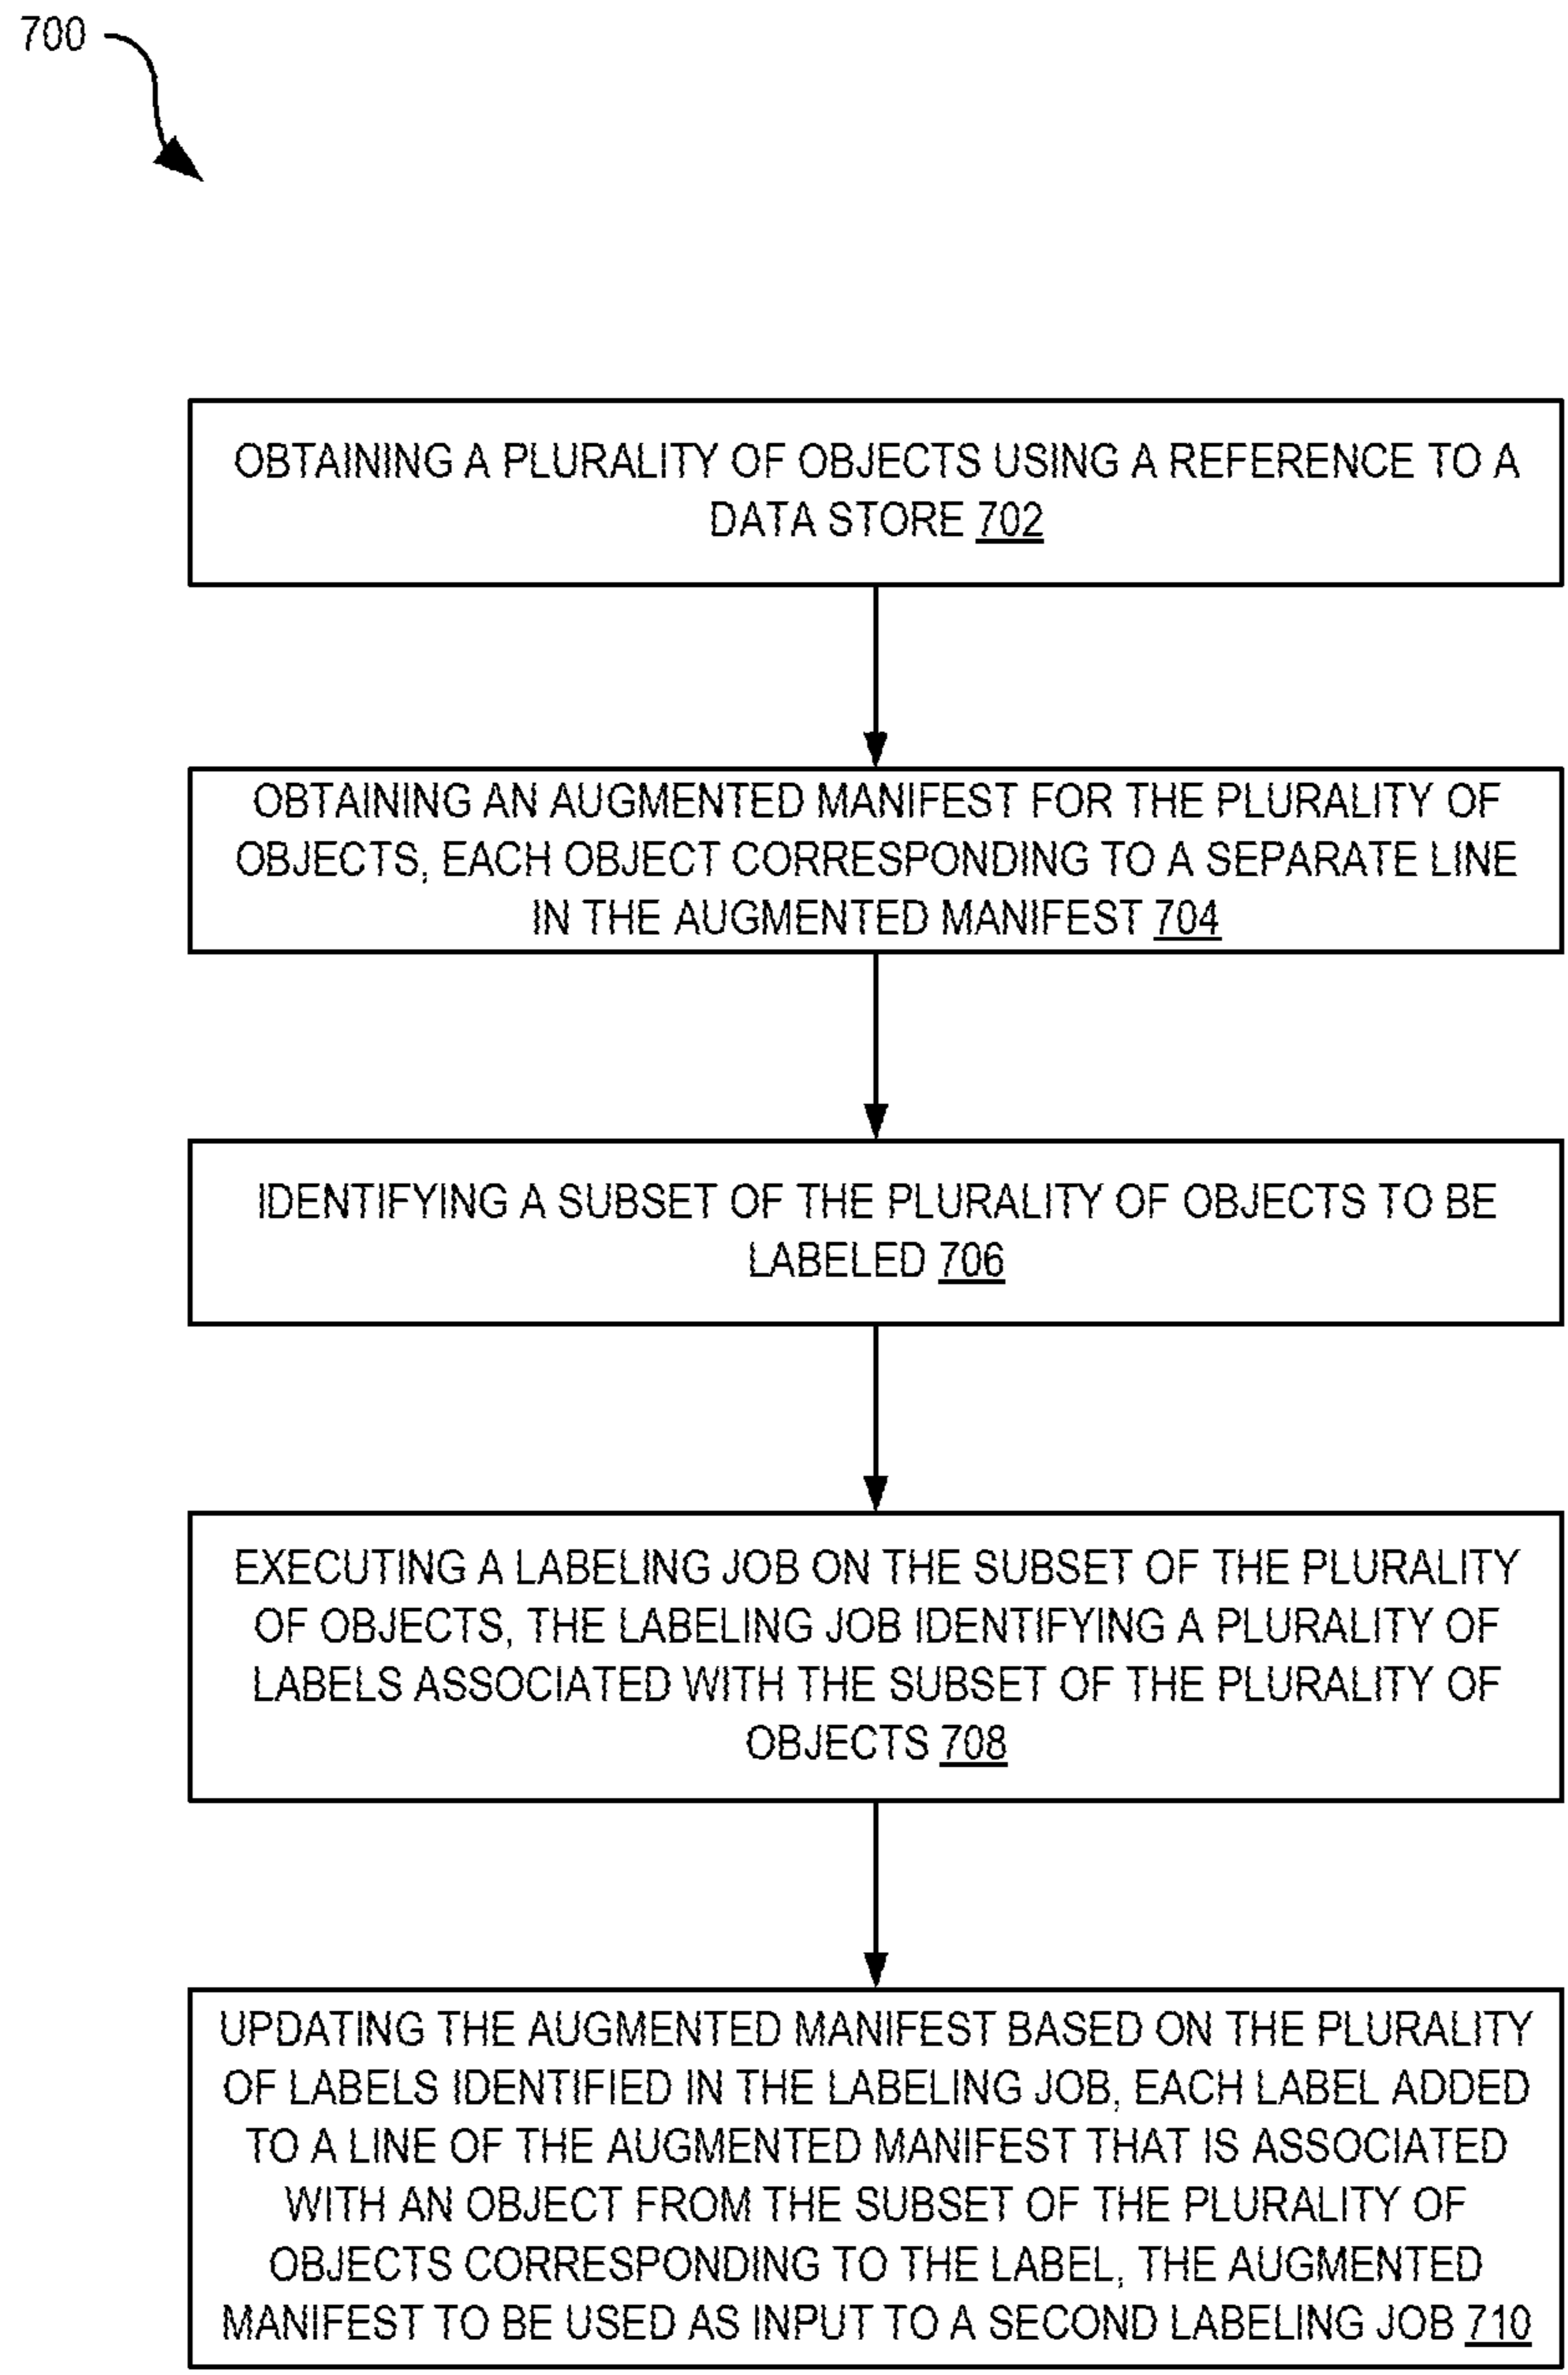


FIG. 7

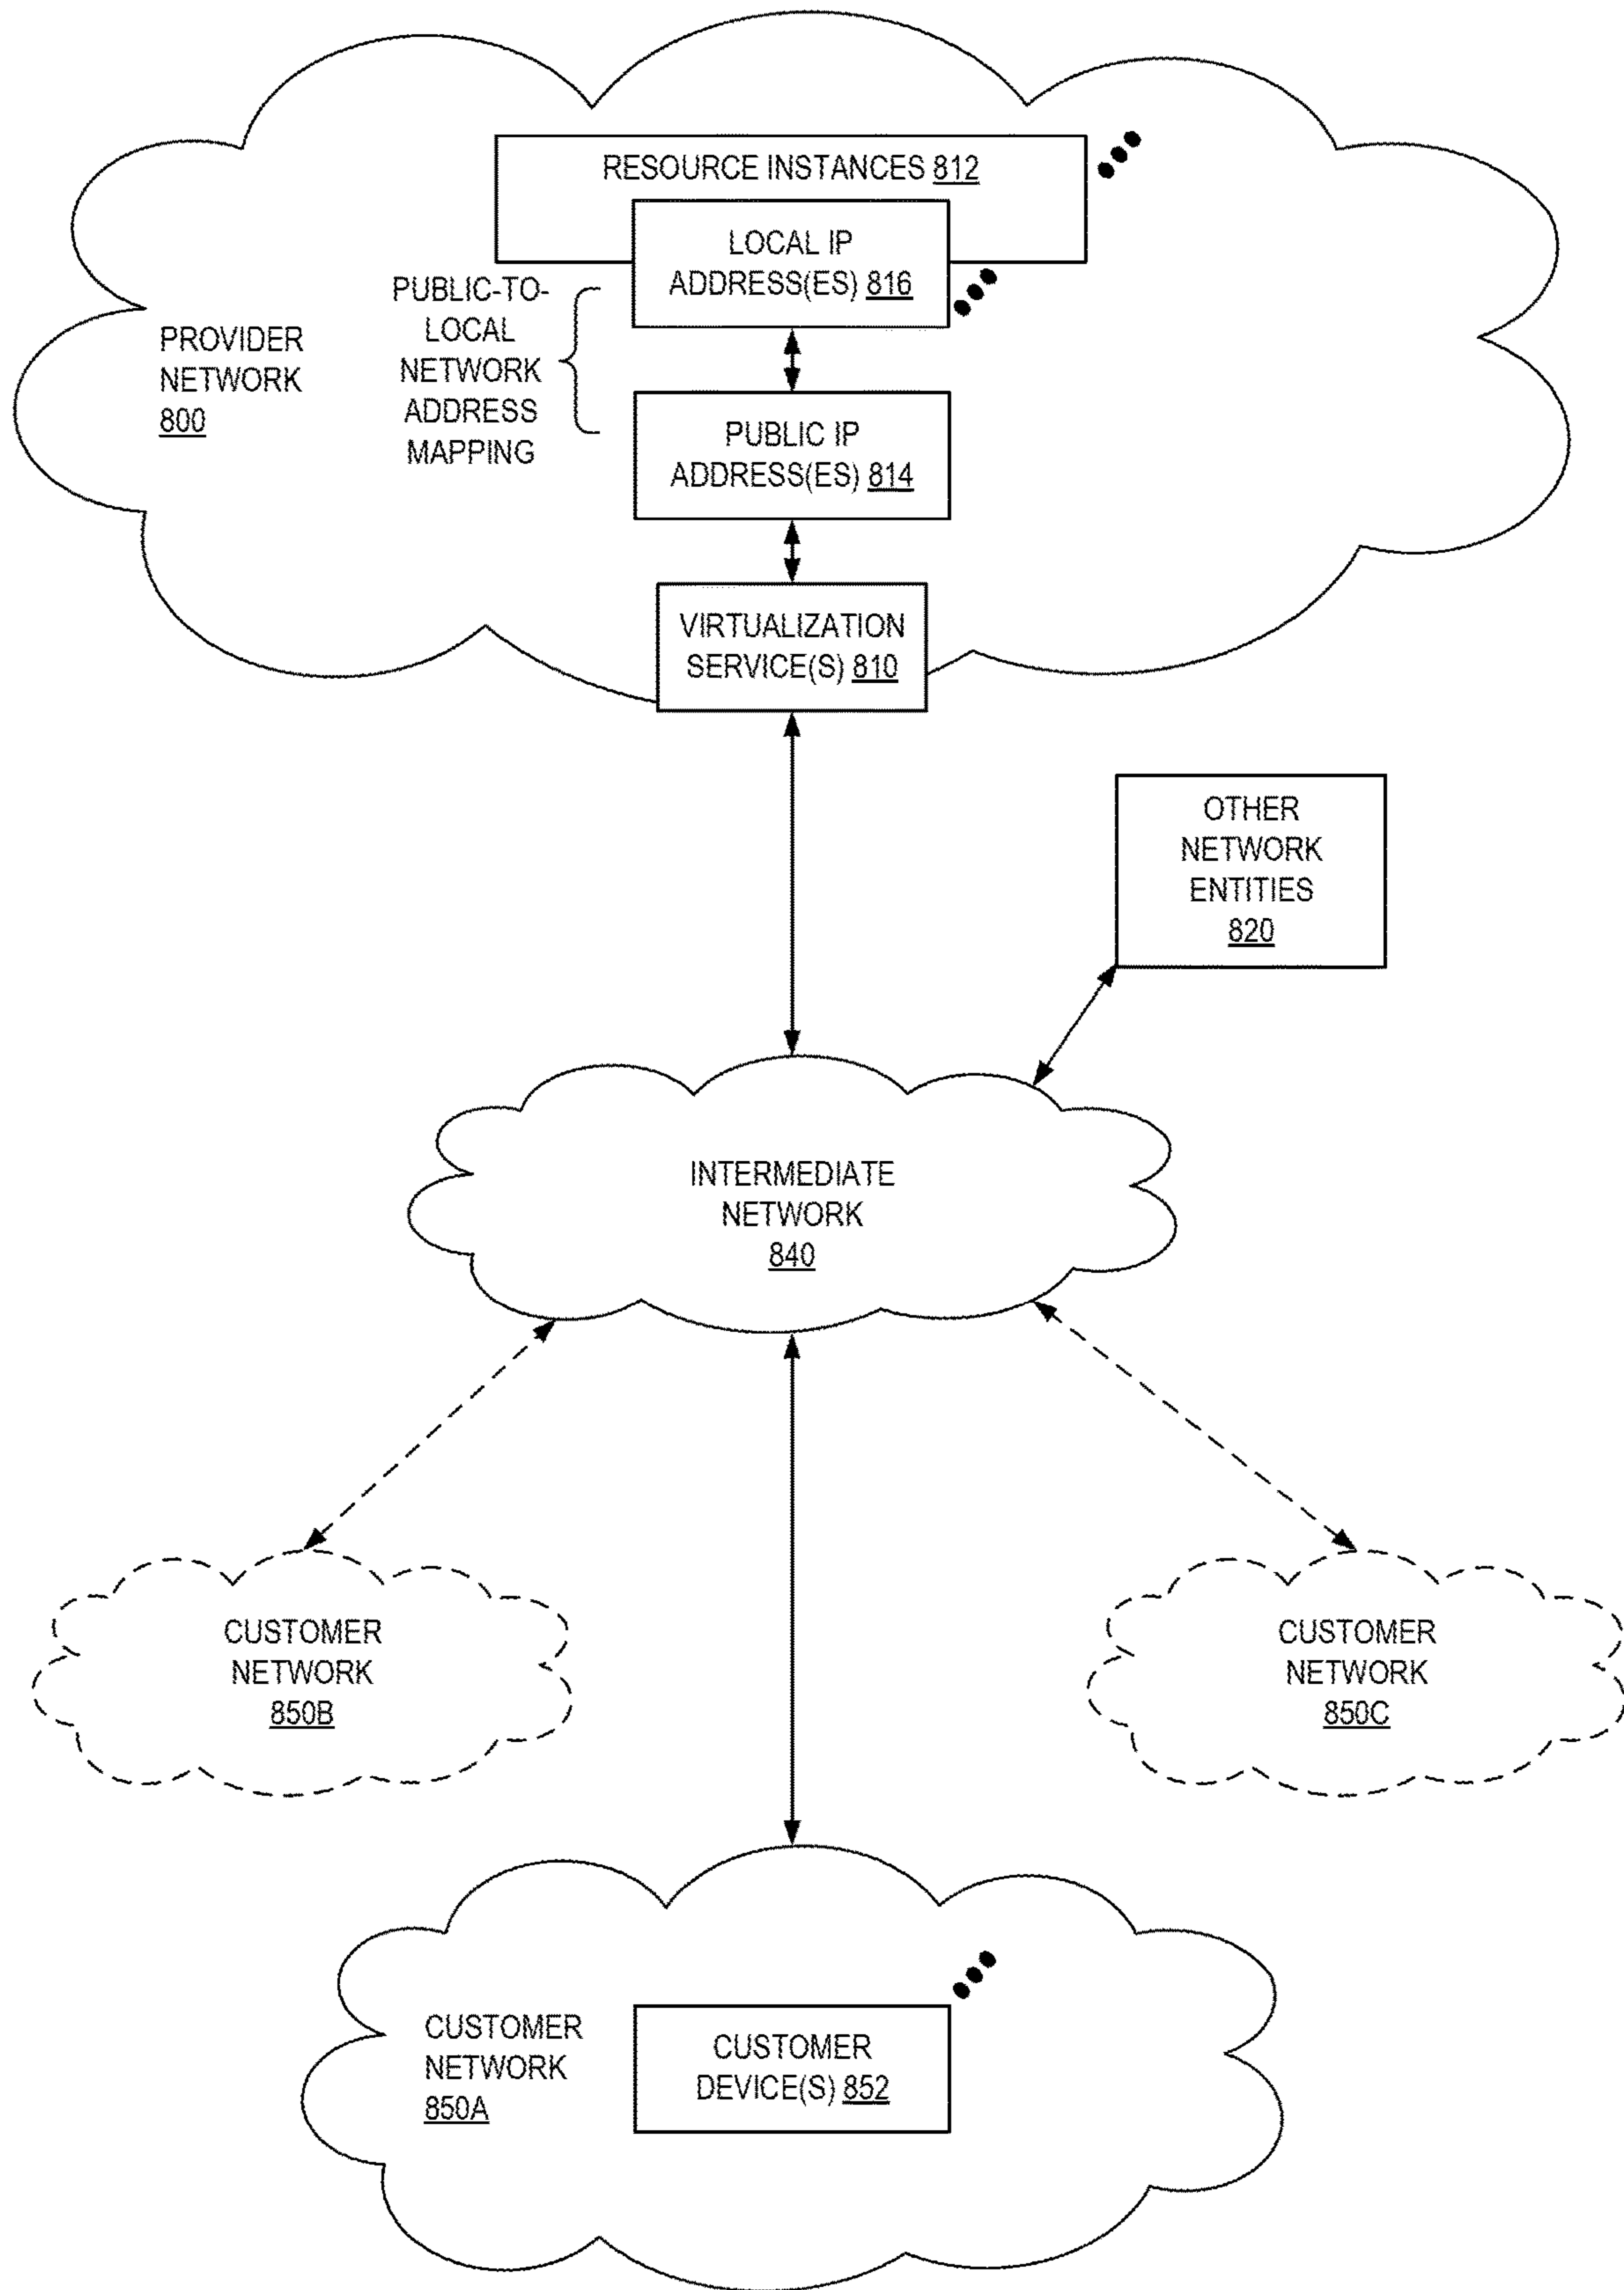


FIG. 8

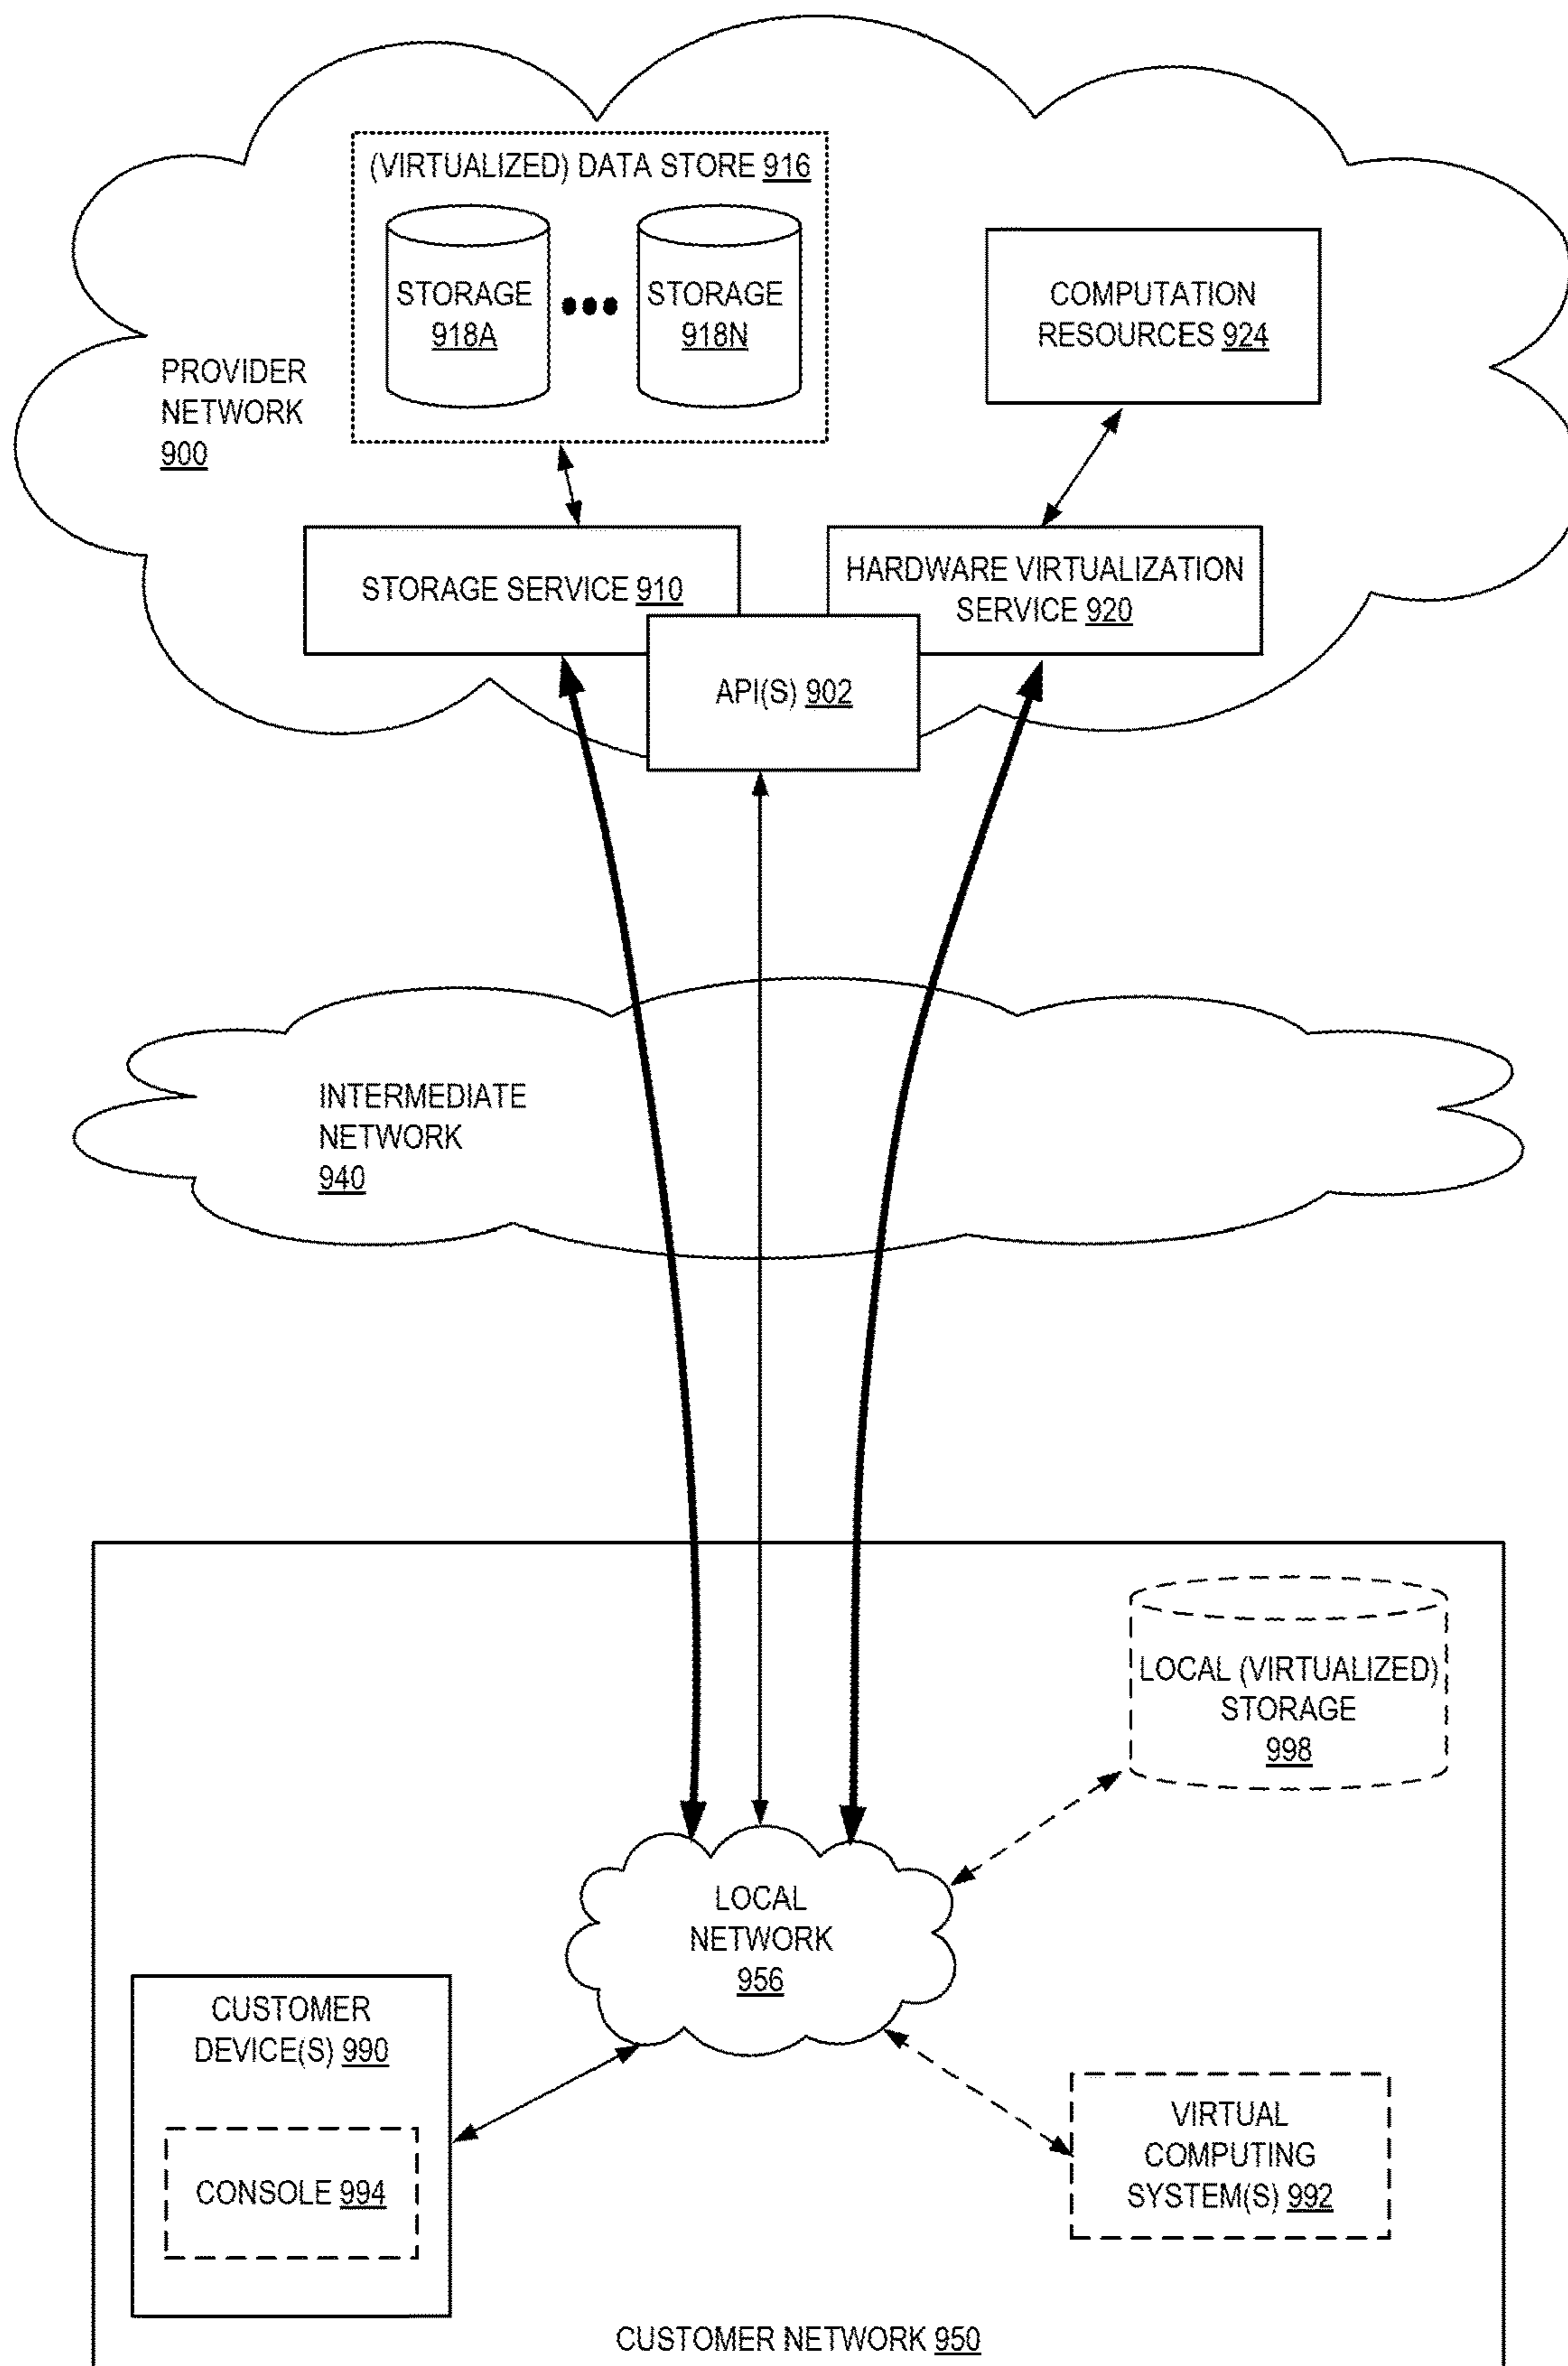


FIG. 9

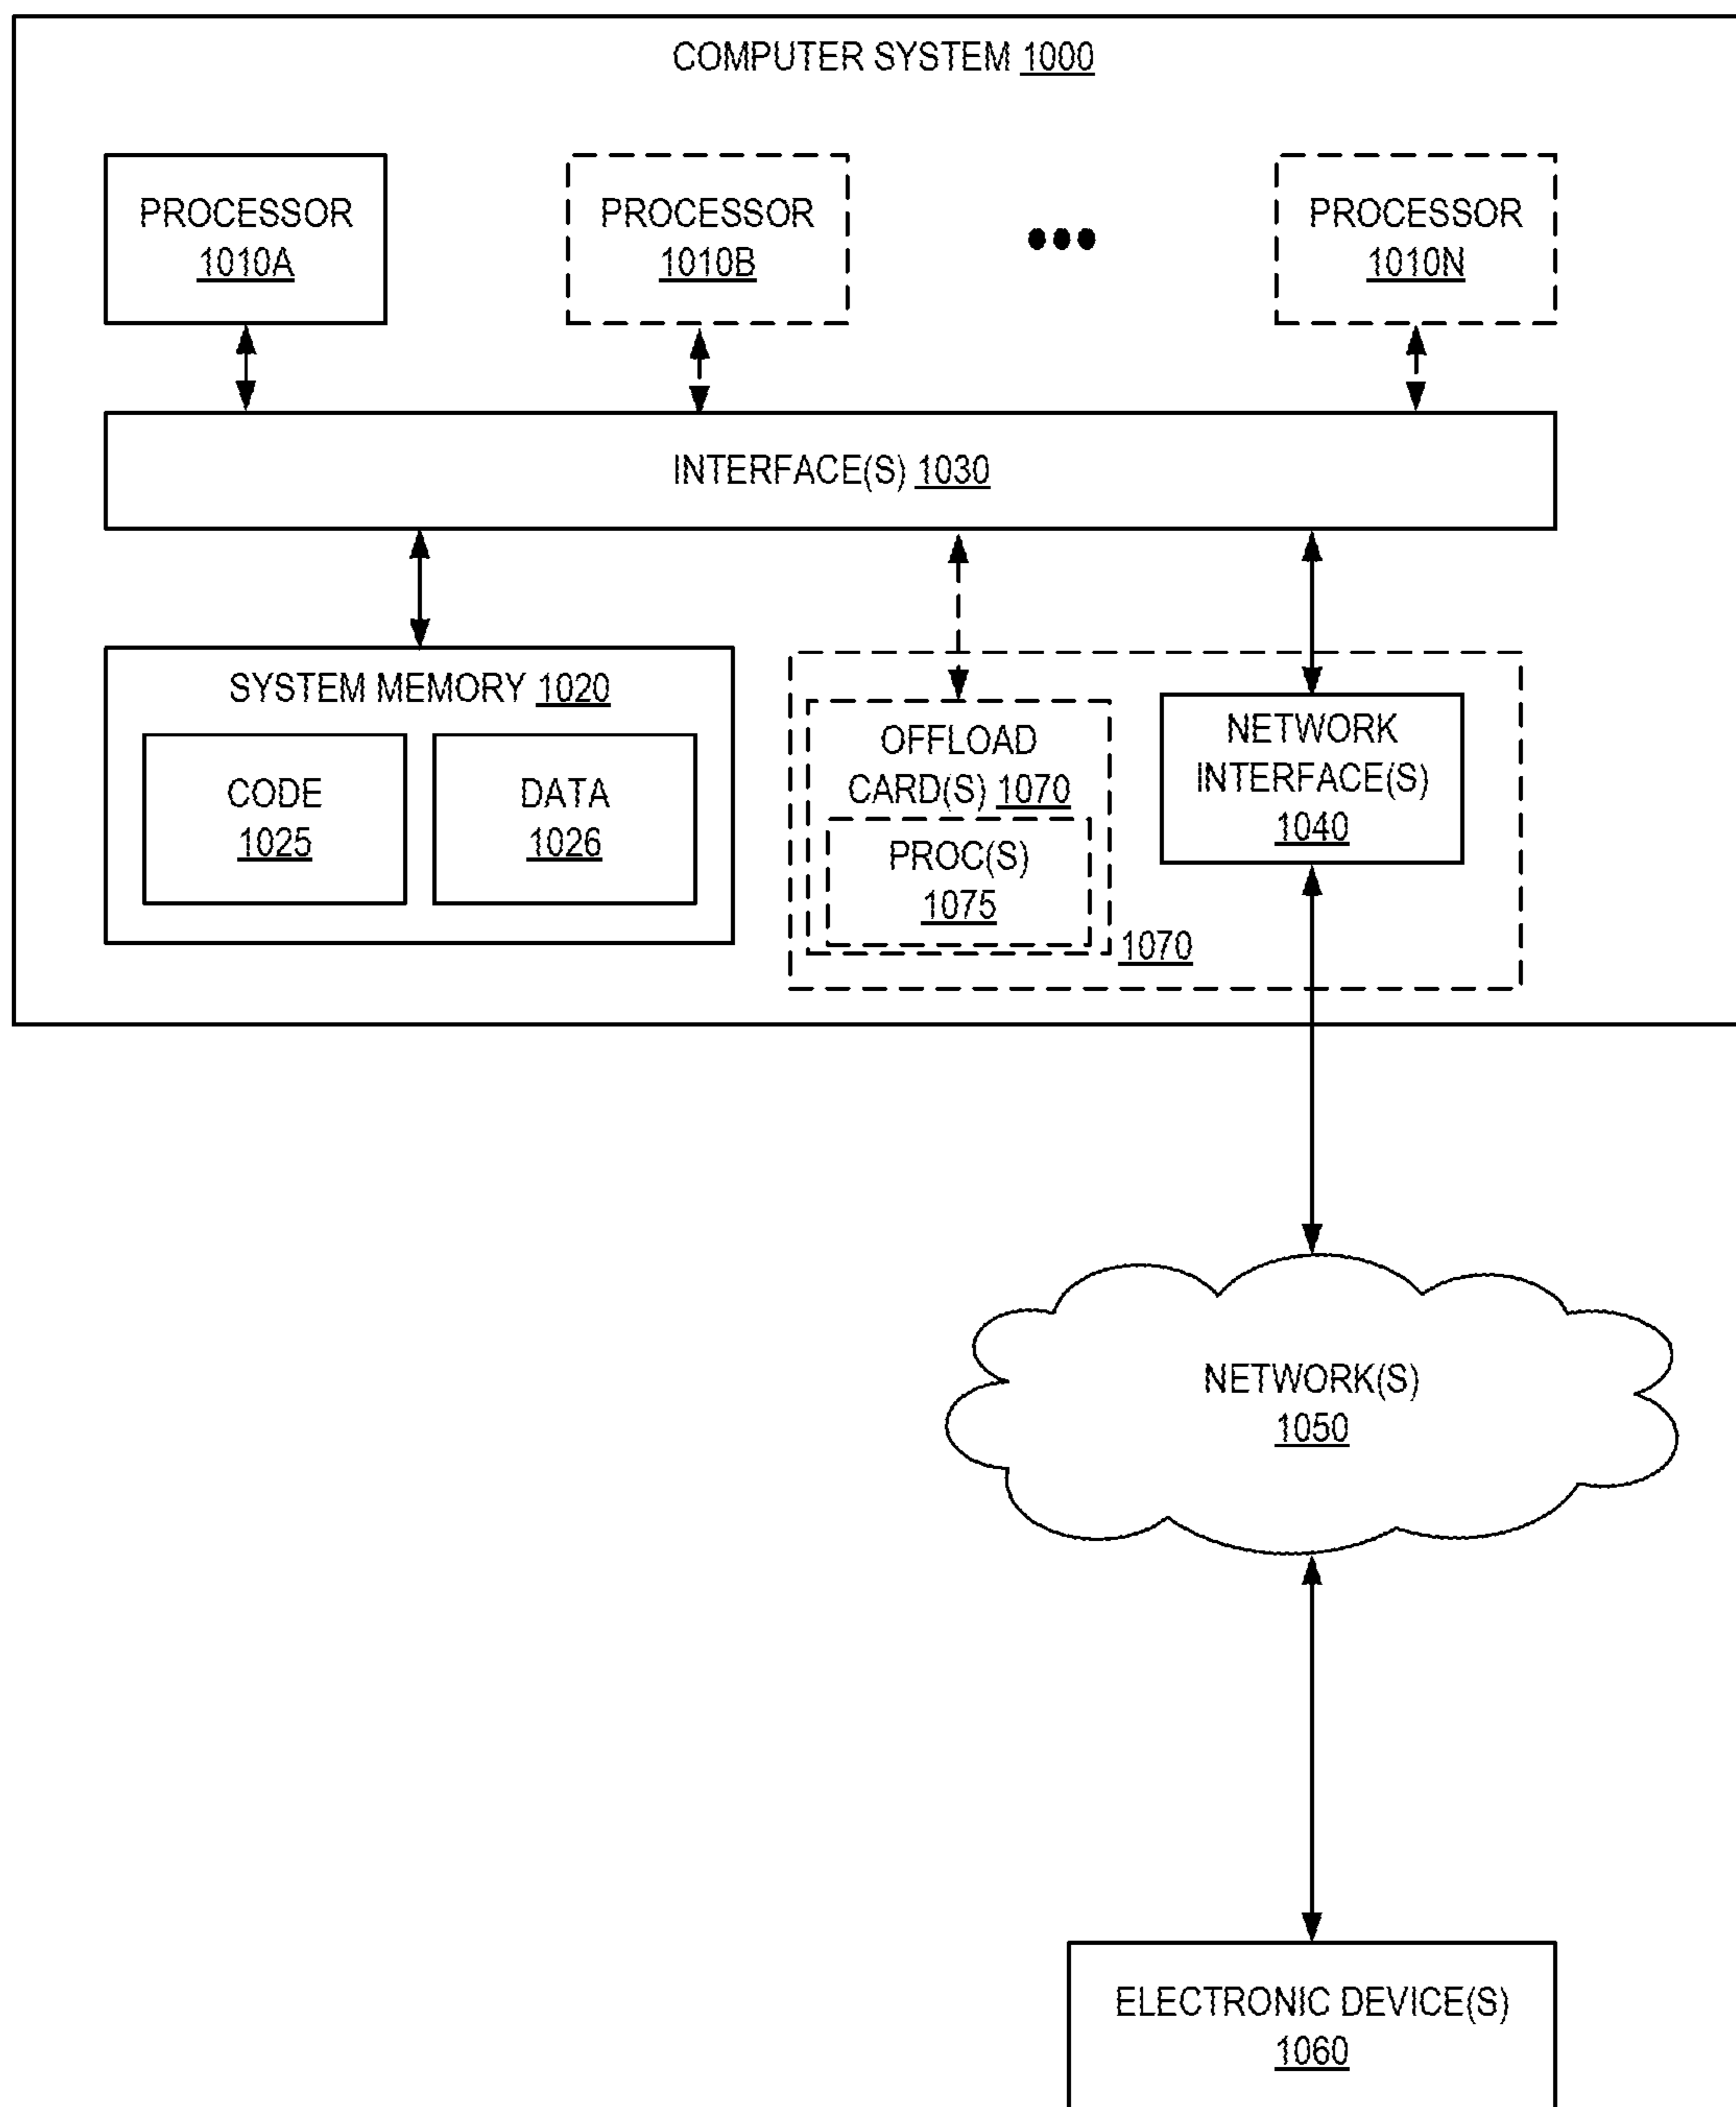


FIG. 10

1

**ACTIVE LEARNING-BASED DATA
LABELING SERVICE USING AN
AUGMENTED MANIFEST**

CROSS-REFERENCE TO RELATED
APPLICATIONS

This application claims the benefit of U.S. Provisional Application No. 62/770,882, filed Nov. 23, 2018, which is hereby incorporated by reference.

BACKGROUND

Many companies and other organizations operate computer networks that interconnect numerous computing systems to support their operations, such as with the computing systems being co-located (e.g., as part of a local network) or instead located in multiple distinct geographical locations (e.g., connected via one or more private or public intermediate networks). For example, data centers housing significant numbers of interconnected computing systems have become commonplace, such as private data centers that are operated by and on behalf of a single organization, and public data centers that are operated by entities as businesses to provide computing resources to customers. Some public data center operators provide network access, power, and secure installation facilities for hardware owned by various customers, while other public data center operators provide “full service” facilities that also include hardware resources made available for use by their customers. However, as the scale and scope of typical data centers has increased, the tasks of provisioning, administering, and managing the physical computing resources have become increasingly complicated.

The advent of virtualization technologies for commodity hardware has provided benefits with respect to managing large-scale computing resources for many customers with diverse needs, allowing various computing resources to be efficiently and securely shared by multiple customers. For example, virtualization technologies may allow a single physical computing machine to be shared among multiple users by providing each user with one or more virtual machines hosted by the single physical computing machine, with each such virtual machine being a software simulation acting as a distinct logical computing system that provides users with the illusion that they are the sole operators and administrators of a given hardware computing resource, while also providing application isolation and security among the various virtual machines. Furthermore, some virtualization technologies are capable of providing virtual resources that span two or more physical resources, such as a single virtual machine with multiple virtual processors that spans multiple distinct physical computing systems. As another example, virtualization technologies may allow data storage hardware to be shared among multiple users by providing each user with a virtualized data store which may be distributed across multiple data storage devices, with each such virtualized data store acting as a distinct logical data store that provides users with the illusion that they are the sole operators and administrators of the data storage resource.

BRIEF DESCRIPTION OF DRAWINGS

Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:

2

FIG. 1 is a diagram illustrating an environment for active learning-based data labeling according to some embodiments.

FIG. 2 is a diagram illustrating an environment for providing data sets to an active learning-based data labeling service and receiving outputs of the active learning-based data labeling service according to some embodiments.

FIG. 3 is a diagram illustrating an environment for active learning-based data labeling using external workforces according to some embodiments.

FIG. 4 is a diagram illustrating labeling of a dataset using an augmented manifest according to some embodiments.

FIG. 5 shows an example augmented manifest according to some embodiments.

FIG. 6 is a diagram illustrating example application programming interfaces (APIs) for active learning-based data labeling according to some embodiments.

FIG. 7 is a flow diagram illustrating operations of a method for active learning-based data labeling using an augmented manifest according to some embodiments.

FIG. 8 illustrates an example provider network environment according to some embodiments.

FIG. 9 is a block diagram of an example provider network that provides a storage service and a hardware virtualization service to customers according to some embodiments.

FIG. 10 is a block diagram illustrating an example computer system that may be used in some embodiments.

DETAILED DESCRIPTION

Various embodiments of methods, apparatus, systems, and non-transitory computer-readable storage media for active learning-based data labeling are described. According to some embodiments, an active learning-based data labeling service enables a user to build and manage large, high accuracy datasets for use in various machine learning systems. In various embodiments, machine learning may be used to automate annotation and management of the datasets, increasing efficiency of labeling tasks and reducing the time required to perform labeling. In the past, labeling has been performed through the brute force efforts of human labelers who are given a dataset to label and return labels (e.g., a description of objects in the dataset, words spoken, bounding boxes around objects, etc.). Embodiments utilize active learning systems to reduce the amount of a dataset that requires manual labeling. As subsets of the dataset are labeled, this label data is used to train a model which can then identify additional objects in the dataset without manual intervention. The process may continue iteratively until the model converges (e.g., identifies objects within an accuracy threshold). This enables a dataset to be labeled without requiring each item in the data set (e.g., image, video frame, video file, audio files, audio tokens, etc.) to be individually and manually labeled by human labelers.

FIG. 1 is a diagram illustrating an environment for active learning-based data labeling according to some embodiments. A provider network **100** provides users with the ability to utilize one or more of a variety of types of computing-related resources such as compute resources (e.g., executing virtual machine (VM) instances and/or containers, executing batch jobs, executing code without provisioning servers), data/storage resources (e.g., object storage, block-level storage, data archival storage, databases and database tables, etc.), network-related resources (e.g., configuring virtual networks including groups of compute resources, content delivery networks (CDNs), Domain Name Service (DNS)), application resources (e.g., data-

bases, application build/deployment services), access policies or roles, identity policies or roles, machine images, routers and other data processing resources, etc. These and other computing resources may be provided as services, such as a hardware virtualization service that can execute compute instances, a storage service that can store data objects, etc. The users (or “customers”) of provider networks **100** may utilize one or more user accounts that are associated with a customer account, though these terms may be used somewhat interchangeably depending upon the context of use. Users may interact with a provider network **100** across one or more intermediate networks **106** (e.g., the internet) via one or more interface(s) **104**, such as through use of application programming interface (API) calls, via a console implemented as a website or application, etc. The interface(s) **104** may be part of, or serve as a front-end to, a control plane of the provider network **100** that includes “backend” services supporting and enabling the services that may be more directly offered to customers.

To provide these and other computing resource services, provider networks **100** often rely upon virtualization techniques. For example, virtualization technologies may be used to provide users the ability to control or utilize compute instances (e.g., a VM using a guest operating system (O/S) that operates using a hypervisor that may or may not further operate on top of an underlying host O/S, a container that may or may not operate in a VM, an instance that can execute on “bare metal” hardware without an underlying hypervisor), where one or multiple compute instances can be implemented using a single electronic device. Thus, a user may directly utilize a compute instance hosted by the provider network to perform a variety of computing tasks, or may indirectly utilize a compute instance by submitting code to be executed by the provider network, which in turn utilizes a compute instance to execute the code (typically without the user having any control of or knowledge of the underlying compute instance(s) involved).

As shown in FIG. 1, at numeral **1**, a customer using electronic device **102** can provide access to a dataset through user interface **104**. As discussed, the user interface **104** can be an API, console, or other interface. For example, the user interface **104** may be a front-end interface for data labeling service **108**. Each dataset can be a collection of homogeneous pieces of data (such as image data, video data, comma separated values (CSV) files, etc.). A dataset may be a raw unlabeled dataset, a partially labeled dataset, a gold standard dataset, or a training dataset. As used herein, a gold standard dataset may refer to a dataset that has been verified as being accurately labeled. In some embodiments, the dataset may be stored in a customer-owned data store, such as an object store, database, or other data store. The customer can provide credentials to access the dataset (e.g., username and password, keys, etc.) and a location of the dataset (e.g., a Uniform Resource Locator (URL) of the dataset or another identifier of a storage location of the dataset). In some embodiments, at numeral **1**, the customer may also specify one or more of a gold standard dataset, a target label space, a desired quality threshold, and an annotation budget.

Labels may be attributes of objects in a dataset. For example, labels may include a region including an object (e.g., a bounding box surrounding a particular object), the species of an animal in a picture, the words in an utterance, etc. Labels are the outputs of annotations after the annotations have been consolidated and have achieved a quality score above a given threshold. As such, as used herein, a label refers to the true underlying object property, while

annotations refer to the tags or other outputs by a labeling task (e.g., by a human labeler or machine annotation).

The label space may be a sub-graph of the overall graph of predefined labels that can be used for annotations for specific dataset objects. For example, for a machine learning model to recognize species of birds, the label space includes labels (and associated label classes) for various species of birds. These labels can be either predefined (e.g., a closed label space) or incrementally defined by annotators (e.g., an open label space)

In some embodiments, the customer can specify a workflow to be used for labeling. The workflow may be provided by the active learning-based labeling service or may be a custom workflow specified by the user. Each workflow can be a nested workflow comprising a series of steps corresponding to one or more of active learning, machine annotating, manual annotating, annotation consolidation and quality score computation, model training, and then outputting of training datasets and/or models.

At numeral **2**, the various properties described above that are received from the customer may be provided to the core engine **110**. The core engine **110** drives the overall workflow definition, execution, monitoring, and orchestration behind the execution of multiple concurrent labeling workflows executed by the data labeling service **108**. The core engine **110** is responsible for triggering/calling functions, activities and tasks on multiple different microservices as part of the overall execution of a workflow. The core engine may also maintain the state (in a multi-tenanted manner) related to the execution of workflows and associated annotation tasks for multiple customers.

At numeral **3**, when the workflow is started, the dataset specified by the customer can be provided to active learning service **112** (“ALS”). ALS **112** may implement one or more active learning techniques as are known in the art. In some embodiments, the active learning technique used may depend on the type of data in the dataset (e.g., image data, audio data, video data, etc.). Additionally, the active learning techniques used may be specified by the customer in a custom workflow or may be built-in as part of data labeling service **108**. The ALS **112** manages the selection of dataset objects that are to be auto-labeled and the selection of dataset objects that are to be manually labeled.

Active learning is a machine learning procedure that can be useful in reducing the amount of annotated data required to achieve a target performance. Active learning starts by incrementally training a model with a small, labeled dataset and then applying this model to the unlabeled data. For each unlabeled sample, ALS **112** estimates whether this sample includes information that has not been learned by the model. An example of an active learning technique is to train an object detection model that takes an image as input and outputs a set of bounding boxes. To train such an object detection model, the training and validation images of the detector are annotated with a bounding box per object and its category. Such a technique may start with a small training set of annotated images to train a baseline object detector. In order to improve the detector by training with more images, this technique continues to collect images to annotate. Rather than annotating all newly collected images, based on different characteristics of the current detector, the ALS **112** can select a subset of the images to be manually labeled. Once annotated, these selected images are added to the training set to incrementally train the object detection model. The entire process continues to collect more images, select a subset with respect to the object detector, annotate the selected ones with humans, incrementally train the

detector and so on. Other data, such as video data, audio data, etc. may also be used in such a system.

ALS 112 can perform active learning for unlabeled or partially unlabeled datasets and use machine learning to evaluate unlabeled raw datasets and provide input into the data labeling process by identifying a subset of the input data to be labeled by manual labelers. In some embodiments, ALS 112 randomly selects a sample of the input dataset for labeling. In some embodiments, ALS 112 selects the subset of the dataset using uncertainty sampling. In this example, a model used by the ALS 112 preferentially selects examples for which the model produces a low confidence of identification. Other approaches may include representativeness-based sampling where the model selects a diverse set that represent the input dataset while limiting redundancy among the subset.

Once a subset of the input dataset is identified to be auto-labeled, the subset may be annotated. For example, in some embodiments, the subset may be sent to machine annotation service 114, as shown at numeral 4. Machine annotation service 114 may use an existing model that has been trained on the same or similar label space which is selected for the input dataset. These machine annotations may be output to a dataset service 116 at numeral 5. In some embodiments, dataset service 116 can communicate with WIS 118, machine annotation service 114, and core engine 110 to manage the flow of data to be labeled by data labeling service 108 and to manage the resulting labels for the data generated by data labeling service 108. Additionally, or alternatively, at numeral 6, the subset identified by the ALS 112 to be manually annotated can be sent to a workforce interface service 118 (“WIS”). The WIS 118 can interface with various manual annotators including crowdsource annotators, private (e.g., in-house) annotators, and/or annotating service providers (collectively annotating service 120). The type and/or variety of labelers may be selected by the customer when starting a new labeling workflow, as discussed above. WIS 118 can present the subset of the input dataset to the selected annotating service 120 at numeral 7. The subset may then be annotated by the selected annotators and the annotations may be returned to WIS 118. These annotations may be stored in dataset service 116 at numeral 8.

Once annotations have been received from the WIS 118 and optionally from machine annotation service 114, the annotations can be consolidated into labels by annotation consolidation service 122, as shown at numeral 9. Annotation consolidation may refer to the process of taking annotations from multiple annotators (e.g., humans and/or machines) and consolidating these together (e.g., using majority-consensus heuristics, removing bias or low-quality annotators, using probabilistic distribution that minimizes a risk function for observed, predicted and true labels, or other techniques). For example, based on each annotators’ accuracy history, their annotations can be weighted. If one annotator has a 50% accurate history, their annotations may have a lower weight than another annotator with a 100% accurate history.

In some embodiments, the annotation consolidation service 122 can maintain a label score and a worker score when performing annotation consolidation. During consolidation, annotation consolidation service 122 can fetch current label scores for each piece of data in the dataset (e.g., image, video frame, audio utterance, etc.) and current worker scores for the annotator who provided the annotations on that piece of data. A new label score and worker scores can be computed for the piece of data based on previous state and

currently received new annotation. The label score can be compared to a specified threshold. If the label score is higher than the threshold then no additional annotations are required. If the label score is lower than the threshold then the data may be passed to additional annotators to be further annotated by annotating service 120. At numeral 10, once the label score is higher than the threshold, then the core engine 110 can be updated to indicate that the subset of the input dataset has been labeled. The active learning loop may continue to execute with the core engine invoking the active learning service 112 to label a new subset of the input dataset that is still unlabeled or partially labeled. The annotation consolidation service can output the labeled subset of the input dataset to an output location, as discussed further below.

In some embodiments, the labeled subset of the input dataset can be used to train the active learning service model. As shown at numeral 11, the labeled subset of the input dataset can be provided to the machine annotation service 114. The machine annotation service 114 may include a training service that can generate a new model, or update the previously used model, using the labeled subset of the input dataset. In some embodiments, a separate training service (not shown) may obtain the labeled subset of the input dataset from the annotation consolidation service and may use the labeled subset of the input dataset to further train the model maintained by the machine annotation service 114. The above described process may then be repeated using the updated model. For example, if the updated model has converged, then the remainder of the input dataset can be accurately identified. If the updated model has not converged, then a new subset of the input dataset can be identified for further labeling according to the process described above. In some embodiments, the data labeling service 108 can output one or more of the converged model or the labeled dataset, as described further below.

In some embodiments, the data labeling service 108 may also output various performance metrics, such as performance against the annotation budget, quality score of annotated labels and performance against the defined quality threshold, logs and metrics in a monitoring dashboard, and/or an audit trail of annotations tasks as performed by annotators. The quality score can be a numerically computed value to measure the quality of consolidated labels. This score is compared to a desired quality threshold (which may be a default value or may be provided by the customer as an input) to determine whether consolidated labels have achieved the desired quality level and hence are considered trained labels. This quality score can be expressed as an output of a risk function on a probability distribution for real, true and predicted labels.

In various embodiments, a customer can provide an unlabeled dataset which they would like to label. They might also bring a pre-trained supervised machine learning model, as well as some data which has already been labeled in the past. These inputs are fed into the data labeling service. The first step of the service is to prepare for the main loop. The preparation produces a validation dataset, a first pre-trained machine learning model, and first estimates of the quality of every worker in the pool. Then the main loop runs, and produces the final model, as well as the fully labeled dataset, with labels for every object of the original unlabeled data, as well as final worker quality estimates.

Before starting the loop, the data labeling service can prepare data for it. The preparation step returns validation data, pretrained model, and worker quality estimate. First, a

random subset of the unlabeled dataset is selected for validation and sent to human annotators. During all of the human annotation processes, worker quality is estimated. There are a few options for the start of the loop. If the customer brings a pre-trained model, it can be passed along to the main loop. If there is no pretrained model, but there is some pre-labeled data, this data is used to train a default model for the selected modality. Even if the labels are not very reliable, pre-labeled data can be used to train the model, as training can often be resilient to some noise in the data. It is not assumed that the pre-labeled data is representative of the unlabeled dataset, and thus the data labeling service still constructs the validation set. Pre-labeled data might have very reliable labels, in which case it can be identified as a Gold Standard dataset, representing true labels and which can be used to help estimate the quality of the human annotations, and individual worker accuracy.

The main loop starts by running inference with the model on the validation dataset. After that, every object is given a confidence level. For example, for image classification, a confidence level can be the probability that the model predicts for the most confident class. A threshold is found, such that for the objects whose model confidence is above the threshold, the expected quality of the label provided by the model exceeds a required quality value provided by the customer. Then, the inference on the unlabeled data is performed, and the threshold is applied on the resulting inferences. All objects with the confidence larger than the threshold get auto-annotated and put into the labeled dataset. For the rest of the objects, active learning is performed to identify a subset of the objects which potentially benefit the model most. Active learning may be executed in a separate active learning loop that iteratively labels portions of the unlabeled dataset and trains the model to better auto-label the dataset, as described herein. This batch of objects is sent for human annotation. During human annotation, previously-found worker quality parameters are used, and adjusted with the new information from their annotations of the new objects. Human labels for these objects are fed into the labeled dataset. All previously human-labeled objects are then used to train a new version of the supervised model, and the loop repeats until all the data is labeled. Then, the labeled dataset is returned to the customer, along with the last model which was trained.

FIG. 2 is a diagram illustrating an environment for providing datasets to an active learning-based data labeling service and receiving outputs of the active learning-based data labeling service according to some embodiments. As shown in FIG. 2, a customer can maintain their own customer data store 200 which includes one or more customer datasets. These may include unlabeled or partially labeled datasets to be provided to data labeling service 108. At numeral 1, the customer (via one or more electronic devices 102) can upload an input dataset to input data store 204. Input data store 204 and output data store 206 may be virtualized data stores 202 provided by a storage service in provider network 100. The storage service may provide object storage, block storage, database storage, or other virtualized storage services. At numeral 2, as discussed above, the customer can provide access information and/or credentials for the input dataset in input data store 204 through a user interface 104. This may include providing a URL for input data store 204 and credentials for accessing the data store. In some embodiments, the customer may establish permissions to allow the data labeling service 108 to access the input dataset. The user interface 104 can pass the information to data labeling service 108 at numeral 3.

At numeral 4, data labeling service 108 can retrieve the input dataset from the input data store 204 and perform active learning-based labeling, as discussed above with respect to FIG. 1. Data labeling service 108 can perform the above described workflow in nested fashion until the active learning model converges. The resulting labeled dataset and/or the converged model can then be output to output data store 206, at numeral 5. Both input and output data stores 204, 206 may be owned by the customer. Any intermediate state data generated for the input dataset by data labeling service 108 can be destroyed, with no durable copies of the labels or input dataset maintained by the data labeling service 108. Thus, the customer's data is labeled and ownership of the dataset and labels is retained by the customer. In some embodiments, the user can access the output data store 206 via user interface 104. For example, the user can send a request to view, download, transfer, etc. the labeled dataset in output data store 206 to user interface 104 at numeral 6. The user interface 104 can obtain the requested labeled dataset or portion thereof at numeral 7 and provide the requested labeled dataset to the user. In some embodiments, the user may access the output data store 206 through a separate user interface (e.g., provided by a storage service that is providing the output data store 206, a local file system, or other interface based on the implementation of the output data store).

In some embodiments, the dataset may include a manifest file which describes dataset properties and records. A record may include named attributes, including metadata such as image size, or labels such as "dog" or "cat". Other attributes may include raw data which needs labeling, such as image or sentences in natural language processing (NLP). In some embodiments, a manifest file for a dataset may be generated automatically by extracting metadata from files in the input data store 204 and generating the manifest file based on the metadata. In some embodiments, the output dataset stored in output data store 206 may be an updated version of the input dataset, which has been updated by combining the input dataset with the generated labels as a new dataset.

FIG. 3 is a diagram illustrating an environment for active learning-based data labeling using external workforces according to some embodiments. As shown in FIG. 3, a customer can create a new labeling job at numeral 1 by sending job details (e.g., input dataset information, credentials, quality threshold, etc.) to user interface 104. At numeral 2, user interface 104 can submit the new labeling job to the core engine 110. In some embodiments, the new labeling job can be submitted to the core engine 110 when the customer submits a start workflow command to the user interface. As described above, the core engine 110 can orchestrate the workflow across various components of data labeling service 108. For example, at numeral 3A, the core engine 110 can provide the input dataset to active learning service 112. Active learning service 112, as described above, can determine a subset of the input dataset to be labeled, either manually or by a machine annotation service. Optionally, in some embodiments, the core engine 110 may provide all or a portion of the input dataset to the workforce interface service as shown at 3B. For example, the active learning service 112 can be bypassed on a first iteration of the active learning loop and a random sample of the input dataset can be provided to the workforce interface service 118 to be labeled. In subsequent iterations, all or a portion of the remaining unlabeled input dataset can be passed to the active learning service 112 to use active learning to increase the speed at which the input dataset is labeled.

At numeral 4, the active learning service 112 can pass the subset identified by the active learning service 112 to be manually annotated to WIS 118. As shown, WIS 118 may include a dataset list 300, which includes each object of the dataset (e.g., image file, text file, video file, video frame, audio utterance, etc.). For each object of the dataset, an annotation manager 302 can create a workflow for annotating service 120 to perform and pass the subset to the annotation service at numeral 5. Annotating service 120, as described above, may enable various workforces to perform annotation of the dataset. For example, annotating service 120 may include an annotating team interface 304, which provides the dataset for annotation to an internal team of annotators, and an external annotating service interface 306, which provides the dataset for annotation to an external annotation team or service such as a third-party annotation vendor. In various embodiments, the dataset may be encrypted prior to it being submitted to an annotation team or service, to protect the dataset while in transit.

The annotated datasets may be returned from the annotators to the annotation manager 302. The annotation manager 302 can output the annotated datasets to dataset service 116 at numeral 6. As discussed above, dataset service 116 can store annotated datasets received from manually annotated and auto-annotated datasets annotated by a machine annotator. At numeral 7, annotation consolidation service 122 can perform annotation consolidation on the annotated datasets. The annotation consolidation service 122 can determine a consolidated annotation (e.g., label) and quality score for each annotated object in the annotated datasets. If the annotation threshold has been reached, the annotation consolidation service 122 can store the resulting labels to output data store 206 at numeral 8.

If the threshold has not been reached, then at numeral 9, the annotation consolidation service can send a request to the annotation manager 302 can extend the annotation workflow execution. This may include requesting a configurable number of additional annotators annotate the dataset. The number of additional annotators may be determined based on the difference between the desired quality threshold and the current threshold. This may continue to loop (e.g., operations depicted as numerals 5-7) for each object of the dataset until all objects of the dataset have been annotated and determined to have a quality score higher than the threshold value.

Augmented Manifest Used for Streaming Datasets for Live Training

FIG. 4 is a diagram illustrating labeling of a dataset using an augmented manifest according to some embodiments. As shown in FIG. 4, a request to start a labeling job can be received through user interface 104. The request can include an augmented data manifest and/or a storage location where data objects are stored to be labeled. In some embodiments, the data to be labeled can be streamed in from a streaming source or can be processed from a static source (e.g., a storage location with a fixed number of data items) in batches. Core engine 110 can instruct dataset service 116 read the user's data to be labeled. If an augmented manifest is received with the request, then there is a fixed dataset size based on the manifest. If the dataset is being streamed, then there is no fixed size, but the dataset service can wait to start a job until a minimum number of objects have been identified (e.g., 10, 20, 1000) depending on user preferences.

When the dataset service has read the user's data, it can output counter values to the core engine 110. The counter values can indicate a total size of the dataset (e.g., X number of objects), a number of labeled objects in the dataset, and

a number of unlabeled objects in the dataset. In the streaming use case, the counters might indicate a number of objects that are ready to be labeled and a number of objects that are still being processed by the dataset service. In some embodiments, the input dataset may be partially labeled. For example, the customer may provide labels for some objects in the dataset with other objects in the dataset being unlabeled. The dataset service 116 can identify which are labeled and which are not labeled based on the manifest during the initial intake. The dataset service can then process only the unlabeled data. This preserves the preexisting labels provided by the customer. If active learning mode is selected by the user, then the core engine can initiate a training mode using the workforce interface service 118.

To begin the labeling job, a batch of data needs to be labeled using the WIS 118. In some embodiments, the WIS does not interface with the dataset service 116 directly, instead it is orchestrated by the core engine. Core engine 110 can send query criteria to the dataset service to obtain the batch to be labeled. The augmented manifest facilitates that query since the metadata for each object are in-line with the objects in the manifest itself. This enables a single query to select the batch of data to be labeled, rather than requiring multiple sources to be queried. In some embodiments, the batch of data can be queried using the query criteria or a random sample of the input dataset can be selected. The dataset service 116 sends the batch to the WIS 118 and then core engine 110 instructs WIS to proceed. As discussed above, WIS 118 can send the batch to annotating service 120 to have the data annotated by one or more types of manual annotators. WIS 118 can then send the resulting annotations to annotation consolidation service 122 to have the annotations consolidated into labels. WIS 118 can provide the labels to the dataset service 116 to update the augmented manifest with the new labels. In some embodiments, the dataset service 116 can maintain a queue to which new labels are added before the augmented manifest is updated (e.g., after the number of new labels enqueued has exceeded a threshold number of labels), or the dataset service can stream in the new labels to the augmented manifest. When the dataset service first reads the manifest file, the dataset service can index the manifest file (e.g., using byte offsets) to identify each line of the manifest file that corresponds to a different object. This enables the dataset service to randomly access the manifest file by a particular line, allowing the dataset service to read from and write to particular lines of the manifest file as needed (e.g., to add new labels to an object, etc.).

Once the initial batch has been labeled, and the augmented manifest updated with the new labels, the dataset service can update the counter values maintained by the core engine 110. Next, active learning can be used on the remaining unlabeled data of the dataset. Once the counter value associated with the number of labeled objects exceeds a threshold value, the core engine can initiate training. The core engine can instruct dataset service 116 to prepare data for training and send it to the active learning service 112. Dataset service 116 can select which data is to be used for training. The training data has to be labeled, needs a confidence score above a threshold, etc. The dataset service can select the training data using the manifest and sends the data to the active learning service. The dataset service can send a message to the core engine indicating that training data has been sent and the core engine can instruct the active learning service to initiate training. In some embodiments, the data sent to the active learning can include training data and a validation dataset. The validation dataset needs to be rep-

representative of the input dataset as a whole. In some embodiments, reservoir sampling, or other large dataset sampling techniques can be used to sample the input dataset to identify a random and representative validation set. The augmented manifest can be updated to add a marker to each object included in the validation dataset. In some embodiments, as the labeling loop iterates, the validation dataset provided initially may no longer be representative of the remaining portion of the dataset to be labeled. The active learning service **112** can send feedback on the validation dataset to the dataset service to remove particular objects from the validation dataset. The dataset service can then unmark the specific items to be removed from the validation dataset. When training is done active learning service **112** can indicate to core engine **110** that it is ready for inference.

Core engine **110** can then instruct dataset service to send unlabeled data to the active learning service. Dataset service **116** sends unlabeled data to the active learning service (or data that is labeled below confidence score, or labeled data that is to be relabeled, etc.). All or a portion of the remaining unlabeled data can then be sent to the active learning service **112**. The active learning service, as discussed, can output two sets of data: labeled data with high confidence scores; and labeled data with low confidence scores. The output datasets are sent back to the dataset service **116** where the labels can be added to the augmented manifest. The data with low confidence score labels can be sampled and sent to WIS **118** for additional manual annotating, and the process can iterate until the input dataset has been completely labeled. The dataset service can continue to update the counters in the core engine as more data is streamed in or made available.

In some embodiments, a user can prepare their data stored in a customer data store and output a dataset of objects with an augmented manifest. The customer data can be stored in a variety of data formats. Such data can include video, image, multi-dimensional vectors, sentences, paragraphs, human-generated labels, and other types of data. The data can include associations between data and labels. In some embodiments, the data format can be CSV, JSON, or others. Large pieces of data, such as images, can be stored as individual objects in a storage service. Many small pieces of data such as sentences in NLP applications may be stored together in one object. The augmented manifest may be organized such that each line of the manifest references a different object in the dataset. Each line can be augmented to include labels from each iteration of the active learning loop. For example, each labeling job generates new labels and metadata for data objects, such as bounding boxes for images. The manifest can be updated to include these new labels and metadata and then used as input for another labeling job and/or another iteration of a current labeling job.

In some embodiments, the labeling working can include the customer performing ETL and/or data preparation on their data in their data store to create a dataset of objects and a manifest. The customer places the manifest file and dataset objects in a storage location in a storage service. The customer can set up permissions to allow data access to the data and manifest. In some embodiments, the customer can provide credentials that provide, e.g., read access to the data and the manifest in the data store. The customer can select or filter their data to create the dataset with the manifest file to be labeled. For example, a query may be executed using the manifest to identify a subset of the dataset to be selected (e.g., using a SELECT statement in SQL, or other filtering condition). The customer starts the labeling job on the

selected dataset. The labeling job stores intermediate results, such as job stats, annotations from human workers, labels, worker states, and dataset object metadata, in a database or other data store. After the labeling job is finished, dataset objects and labels can be combined into a new output dataset in an output storage location associated with the customer. The intermediate job states are deleted when the labeling job finishes. The customer can use the output dataset for another labeling job, prep/transform dataset for training a machine learning model, or store labels to any data store the customer owns.

FIG. **5** shows an example augmented manifest according to some embodiments. A record **500** includes named attributes. Some attributes are metadata such as image size, or labels such as “dog” “cat”. Other attributes are raw data which needs labeling, such as an image or sentences in NLP. As shown in FIG. **5**, attributes can be large, represented as references to objects stored in a storage service, or small, represented inline. In some embodiments, attributes can be complex JSON objects.

In some embodiments, a user can create a dataset without using any ETL tool to create a manifest first. A workflow can crawl over a folder in the customer’s data store and look for JPEG files, text files, or other file formats corresponding to a supported data type. It extracts metadata automatically from the files which conforms to a specific format. It generates a dataset manifest in the customer’s data store at a specified location.

In some embodiments, at the end of a labeling job, the manifest can be updated by combining the input dataset and labels generated by the job as a new dataset. The new dataset (with new manifest file) is stored at a location specified in the customer’s data store. In some embodiments, the manifest of a new dataset may include only the delta (e.g., diff) and a reference to the old dataset. In some embodiments, labels of a dataset can be edited in the console.

In some embodiments, the dataset can be filtered using queries on the augmented manifest file using common data types such as string, number, dates, JSON. In some embodiments, the queries may be filtering expressions similar to WHERE clause in SQL.

Embodiments can track data lineage by tracking change history of dataset and relationship between trained models and dataset. Embodiments can also run a labeling job on a random sampling of dataset. Customers may provide parameters to random sampling to have distinct sub-groups of data samples. In some embodiments, labels may be hierarchical (for example, person: boy, girl).

The manifest enables datasets to be streamed into the active learning pipeline for labeling and live training. The manifest does not require the user to separate out the data source and label in separate files or data structures, as the manifest can be augmented to include the labels in-line with the data or a reference to the data. The dataset is also designed so that the user does not have to define where records begin or end using a specific preamble or post-script. Instead, each record corresponds to one line of the manifest, without requiring explicitly declared boundaries. For example, a dataset may have been labeled with “object” or “no object” indicating that a particular object either has been detected or not detected in an image. The dataset may be filtered to run an additional label job (e.g., to add a bounding box, add additional labels, etc.) on just those pieces of data in the dataset that include the label “object”.

FIG. **6** is a diagram illustrating example application programming interfaces (APIs) **600** for active learning-based data labeling according to some embodiments. As

shown in FIG. 6, the example APIs **600** may include a CreateLabelingJob **602** command which may receive a dataset URL or other identifier, a dataset type (e.g., raw, partially labeled, gold standard, etc.), a quality threshold, and an annotation budget. The CreateLabelingJob command can cause a new labeling project to be created and coordinated by the core engine of the training dataset management service. The API **600** may further include a DescribeLabelingJob command **604** which may receive a customer ID and return projects associated with that customer ID. The projects may include currently pending projects and/or completed projects. The API **600** may further include a ListLabelingJobsForWorkTeam command **606** which may return all labeling jobs associated with a workteam identifier. In some embodiments, the API **600** may further include a ListLabelingJobs command **608** that may return objects representing current labeling jobs (e.g., pending labeling jobs). In some embodiments, the API **600** may further include a StopLabelingJob command **610** that may receive an identifier associated with a labeling job and stops execution of the labeling job. In some embodiments, the API **600** may also include a start labeling job command that may cause the core engine to start a labeling workflow for a given project.

In some embodiments, the API **600** may also include various other commands, such as a GetWorkflowOutput command, which can receive a workflow ID and return the output of the workflow including a training dataset and labels (e.g., the labeled input dataset for that workflow) or a trained model. The API **600** may also include a GetLabelsForDatasetObject command which receives an object ID and returns all labels associated with that object. For example, an image file ID (e.g., a file name or other identifier) may be received by this command and all labels for that image may be returned. Similarly, a GetDatasetObjectsforLabel command can receive a label and return a list of all objects in the dataset that are associated with that specific label.

FIG. 7 is a flow diagram illustrating operations of a method **700** for active learning-based data labeling according to some embodiments. Some or all of the operations **700** (or other processes described herein, or variations, and/or combinations thereof) are performed under the control of one or more computer systems configured with executable instructions and are implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware or combinations thereof. The code is stored on a computer-readable storage medium, for example, in the form of a computer program comprising instructions executable by one or more processors. The computer-readable storage medium is non-transitory. In some embodiments, one or more (or all) of the operations **700** are performed by data labeling service **108** of the other figures.

The operations **700** include, at block **702**, obtaining a plurality of objects using a reference to a data store. In some embodiments, receiving a reference can include receiving credentials to read the plurality of objects and the augmented manifest. The operations **700** include, at block **704**, obtaining an augmented manifest for the plurality of objects, each object corresponding to a separate line in the augmented manifest. In some embodiments, the plurality of objects include a reference to at least one of a video, an image, a multi-dimensional vector, text data, or a human-generated label. In some embodiments, obtaining an augmented manifest for the plurality of objects, each object corresponding to a separate line in the augmented manifest further comprises:

reading the augmented manifest from the data store, and indexing each line of the augmented manifest.

The operations **700** include, at block **706**, identifying a subset of the plurality of objects to be labeled. In some embodiments, obtaining the augmented manifest may include identifying a plurality of files of a data type in the data store, the plurality of files associated with the plurality of objects, extracting metadata from the plurality of files, and generating the augmented manifest in the data store based on the metadata. In some embodiments, identifying a subset of the plurality of objects to be labeled further comprises executing a query on the augmented manifest to identify the subset, wherein the query includes a filtering condition executed on the augmented manifest to filter the plurality of objects. In some embodiments, identifying a subset of the plurality of objects to be labeled further comprises randomly sampling the plurality of objects using a large dataset sampling technique.

The operations **700** include, at block **708**, executing a labeling job on the subset of the plurality of objects, the labeling job identifying a plurality of labels associated with the subset of the plurality of objects. In some embodiments, the query includes a select statement executed on the augmented manifest to filter the plurality of objects.

The operations **700** include, at block **710**, updating the augmented manifest based on the plurality of labels identified in the labeling job, each label added to a line of the augmented manifest that is associated with an object from the subset of the plurality of objects corresponding to the label, the augmented manifest to be used as input to a second labeling job. In some embodiments, each object is defined in the augmented manifest without using a preamble or postscript to define boundaries of the object in the augmented manifest. In some embodiments, updating the augmented manifest based on the plurality of labels identified in the labeling job, further comprises identifying a line of the augmented manifest corresponding to a first label from the plurality of labels, and writing the first label to the line of the augmented manifest using the byte offset indexed for that line.

In some embodiments, the operations may further include storing intermediate results of the labeling job, the intermediate results including at least one of job statistics, annotations obtained during the labeling job, labels, worker states, or object metadata. In some embodiments, the operations may further include combining the plurality of objects and the plurality of labels into a new output dataset, and storing the new output dataset and a new augmented manifest in an output data store. In some embodiments, the new augmented manifest includes changes corresponding to the labeling job and a reference to the plurality of objects.

In some embodiments, the operations may include receiving a reference to a data store and credentials to access the data store, identifying text data in the data store, the text data including a plurality of objects, obtaining an augmented manifest for the text data, each object from the plurality of objects corresponding to a separate line in the augmented manifest, executing a labeling job on a subset of the plurality of objects, the subset of the plurality of objects identified by running a query on the augmented manifest, the labeling job identifying a plurality of labels associated with the subset of the plurality of objects, adding the plurality of labels to the augmented manifest to generate a new augmented manifest, each label added in-line with its corresponding object, the augmented manifest to be used as input to a second labeling job, and storing an output dataset and the new augmented manifest in an output data store.

FIG. 8 illustrates an example provider network (or “service provider system”) environment according to some embodiments. A provider network **800** may provide resource virtualization to customers via one or more virtualization services **810** that allow customers to purchase, rent, or otherwise obtain instances **812** of virtualized resources, including but not limited to computation and storage resources, implemented on devices within the provider network or networks in one or more data centers. Local Internet Protocol (IP) addresses **816** may be associated with the resource instances **812**; the local IP addresses are the internal network addresses of the resource instances **812** on the provider network **800**. In some embodiments, the provider network **800** may also provide public IP addresses **814** and/or public IP address ranges (e.g., Internet Protocol version 4 (IPv4) or Internet Protocol version 8 (IPv6) addresses) that customers may obtain from the provider **800**.

Conventionally, the provider network **800**, via the virtualization services **810**, may allow a customer of the service provider (e.g., a customer that operates one or more client networks **850A-850C** including one or more customer device(s) **852**) to dynamically associate at least some public IP addresses **814** assigned or allocated to the customer with particular resource instances **812** assigned to the customer. The provider network **800** may also allow the customer to remap a public IP address **814**, previously mapped to one virtualized computing resource instance **812** allocated to the customer, to another virtualized computing resource instance **812** that is also allocated to the customer. Using the virtualized computing resource instances **812** and public IP addresses **814** provided by the service provider, a customer of the service provider such as the operator of customer network(s) **850A-850C** may, for example, implement customer-specific applications and present the customer’s applications on an intermediate network **840**, such as the Internet. Other network entities **820** on the intermediate network **840** may then generate traffic to a destination public IP address **814** published by the customer network(s) **850A-850C**; the traffic is routed to the service provider data center, and at the data center is routed, via a network substrate, to the local IP address **816** of the virtualized computing resource instance **812** currently mapped to the destination public IP address **814**. Similarly, response traffic from the virtualized computing resource instance **812** may be routed via the network substrate back onto the intermediate network **840** to the source entity **820**.

Local IP addresses, as used herein, refer to the internal or “private” network addresses, for example, of resource instances in a provider network. Local IP addresses can be within address blocks reserved by Internet Engineering Task Force (IETF) Request for Comments (RFC) 1918 and/or of an address format specified by IETF RFC 4193, and may be mutable within the provider network. Network traffic originating outside the provider network is not directly routed to local IP addresses; instead, the traffic uses public IP addresses that are mapped to the local IP addresses of the resource instances. The provider network may include networking devices or appliances that provide network address translation (NAT) or similar functionality to perform the mapping from public IP addresses to local IP addresses and vice versa.

Public IP addresses are Internet mutable network addresses that are assigned to resource instances, either by the service provider or by the customer. Traffic routed to a public IP address is translated, for example via 1:1 NAT, and forwarded to the respective local IP address of a resource instance.

Some public IP addresses may be assigned by the provider network infrastructure to particular resource instances; these public IP addresses may be referred to as standard public IP addresses, or simply standard IP addresses. In some embodiments, the mapping of a standard IP address to a local IP address of a resource instance is the default launch configuration for all resource instance types.

At least some public IP addresses may be allocated to or obtained by customers of the provider network **800**; a customer may then assign their allocated public IP addresses to particular resource instances allocated to the customer. These public IP addresses may be referred to as customer public IP addresses, or simply customer IP addresses. Instead of being assigned by the provider network **800** to resource instances as in the case of standard IP addresses, customer IP addresses may be assigned to resource instances by the customers, for example via an API provided by the service provider. Unlike standard IP addresses, customer IP addresses are allocated to customer accounts and can be remapped to other resource instances by the respective customers as necessary or desired. A customer IP address is associated with a customer’s account, not a particular resource instance, and the customer controls that IP address until the customer chooses to release it. Unlike conventional static IP addresses, customer IP addresses allow the customer to mask resource instance or availability zone failures by remapping the customer’s public IP addresses to any resource instance associated with the customer’s account. The customer IP addresses, for example, enable a customer to engineer around problems with the customer’s resource instances or software by remapping customer IP addresses to replacement resource instances.

FIG. 9 is a block diagram of an example provider network that provides a storage service and a hardware virtualization service to customers, according to some embodiments. Hardware virtualization service **920** provides multiple computation resources **924** (e.g., VMs) to customers. The computation resources **924** may, for example, be rented or leased to customers of the provider network **900** (e.g., to a customer that implements customer network **950**). Each computation resource **924** may be provided with one or more local IP addresses. Provider network **900** may be configured to route packets from the local IP addresses of the computation resources **924** to public Internet destinations, and from public Internet sources to the local IP addresses of computation resources **924**.

Provider network **900** may provide a customer network **950**, for example coupled to intermediate network **940** via local network **956**, the ability to implement virtual computing systems **992** via hardware virtualization service **920** coupled to intermediate network **940** and to provider network **900**. In some embodiments, hardware virtualization service **920** may provide one or more APIs **902**, for example a web services interface, via which a customer network **950** may access functionality provided by the hardware virtualization service **920**, for example via a console **994** (e.g., a web-based application, standalone application, mobile application, etc.). In some embodiments, at the provider network **900**, each virtual computing system **992** at customer network **950** may correspond to a computation resource **924** that is leased, rented, or otherwise provided to customer network **950**.

From an instance of a virtual computing system **992** and/or another customer device **990** (e.g., via console **994**), the customer may access the functionality of storage service **910**, for example via one or more APIs **902**, to access data from and store data to storage resources **918A-918N** of a

virtual data store **916** (e.g., a folder or “bucket”, a virtualized volume, a database, etc.) provided by the provider network **900**. In some embodiments, a virtualized data store gateway (not shown) may be provided at the customer network **950** that may locally cache at least some data, for example frequently-accessed or critical data, and that may communicate with storage service **910** via one or more communications channels to upload new or modified data from a local cache so that the primary store of data (virtualized data store **916**) is maintained. In some embodiments, a user, via a virtual computing system **992** and/or on another customer device **990**, may mount and access virtual data store **916** volumes via storage service **910** acting as a storage virtualization service, and these volumes may appear to the user as local (virtualized) storage **998**.

While not shown in FIG. 9, the virtualization service(s) may also be accessed from resource instances within the provider network **900** via API(s) **902**. For example, a customer, appliance service provider, or other entity may access a virtualization service from within a respective virtual network on the provider network **900** via an API **902** to request allocation of one or more resource instances within the virtual network or within another virtual network. Illustrative System

In some embodiments, a system that implements a portion or all of the techniques for active learning-based data labeling as described herein may include a general-purpose computer system that includes or is configured to access one or more computer-accessible media, such as computer system **1000** illustrated in FIG. 10. In the illustrated embodiment, computer system **1000** includes one or more processors **1010** coupled to a system memory **1020** via an input/output (I/O) interface **1030**. Computer system **1000** further includes a network interface **1040** coupled to I/O interface **1030**. While FIG. 10 shows computer system **1000** as a single computing device, in various embodiments a computer system **1000** may include one computing device or any number of computing devices configured to work together as a single computer system **1000**.

In various embodiments, computer system **1000** may be a uniprocessor system including one processor **1010**, or a multiprocessor system including several processors **1010** (e.g., two, four, eight, or another suitable number). Processors **1010** may be any suitable processors capable of executing instructions. For example, in various embodiments, processors **1010** may be general-purpose or embedded processors implementing any of a variety of instruction set architectures (ISAs), such as the x86, ARM, PowerPC, SPARC, or MIPS ISAs, or any other suitable ISA. In multiprocessor systems, each of processors **1010** may commonly, but not necessarily, implement the same ISA.

System memory **1020** may store instructions and data accessible by processor(s) **1010**. In various embodiments, system memory **1020** may be implemented using any suitable memory technology, such as random-access memory (RAM), static RAM (SRAM), synchronous dynamic RAM (SDRAM), nonvolatile/Flash-type memory, or any other type of memory. In the illustrated embodiment, program instructions and data implementing one or more desired functions, such as those methods, techniques, and data described above are shown stored within system memory **1020** as code **1025** and data **1026**.

In one embodiment, I/O interface **1030** may be configured to coordinate I/O traffic between processor **1010**, system memory **1020**, and any peripheral devices in the device, including network interface **1040** or other peripheral interfaces. In some embodiments, I/O interface **1030** may per-

form any necessary protocol, timing or other data transformations to convert data signals from one component (e.g., system memory **1020**) into a format suitable for use by another component (e.g., processor **1010**). In some embodiments, I/O interface **1030** may include support for devices attached through various types of peripheral buses, such as a variant of the Peripheral Component Interconnect (PCI) bus standard or the Universal Serial Bus (USB) standard, for example. In some embodiments, the function of I/O interface **1030** may be split into two or more separate components, such as a north bridge and a south bridge, for example. Also, in some embodiments some or all of the functionality of I/O interface **1030**, such as an interface to system memory **1020**, may be incorporated directly into processor **1010**.

Network interface **1040** may be configured to allow data to be exchanged between computer system **1000** and other devices **1060** attached to a network or networks **1050**, such as other computer systems or devices as illustrated in FIG. 1, for example. In various embodiments, network interface **1040** may support communication via any suitable wired or wireless general data networks, such as types of Ethernet network, for example. Additionally, network interface **1040** may support communication via telecommunications/telephony networks such as analog voice networks or digital fiber communications networks, via storage area networks (SANs) such as Fibre Channel SANs, or via I/O any other suitable type of network and/or protocol.

In some embodiments, a computer system **1000** includes one or more offload cards **1070** (including one or more processors **1075**, and possibly including the one or more network interfaces **1040**) that are connected using an I/O interface **1030** (e.g., a bus implementing a version of the Peripheral Component Interconnect-Express (PCI-E) standard, or another interconnect such as a QuickPath interconnect (QPI) or UltraPath interconnect (UPI)). For example, in some embodiments the computer system **1000** may act as a host electronic device (e.g., operating as part of a hardware virtualization service) that hosts compute instances, and the one or more offload cards **1070** execute a virtualization manager that can manage compute instances that execute on the host electronic device. As an example, in some embodiments the offload card(s) **1070** can perform compute instance management operations such as pausing and/or un-pausing compute instances, launching and/or terminating compute instances, performing memory transfer/copying operations, etc. These management operations may, in some embodiments, be performed by the offload card(s) **1070** in coordination with a hypervisor (e.g., upon a request from a hypervisor) that is executed by the other processors **1010A-1010N** of the computer system **1000**. However, in some embodiments the virtualization manager implemented by the offload card(s) **1070** can accommodate requests from other entities (e.g., from compute instances themselves), and may not coordinate with (or service) any separate hypervisor.

In some embodiments, system memory **1020** may be one embodiment of a computer-accessible medium configured to store program instructions and data as described above. However, in other embodiments, program instructions and/or data may be received, sent or stored upon different types of computer-accessible media. Generally speaking, a computer-accessible medium may include non-transitory storage media or memory media such as magnetic or optical media, e.g., disk or DVD/CD coupled to computer system **1000** via I/O interface **1030**. A non-transitory computer-accessible storage medium may also include any volatile or non-volatile media such as RAM (e.g., SDRAM, double data rate

(DDR) SDRAM, SRAM, etc.), read only memory (ROM), etc., that may be included in some embodiments of computer system **1000** as system memory **1020** or another type of memory. Further, a computer-accessible medium may include transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as a network and/or a wireless link, such as may be implemented via network interface **1040**.

Various embodiments discussed or suggested herein can be implemented in a wide variety of operating environments, which in some cases can include one or more user computers, computing devices, or processing devices which can be used to operate any of a number of applications. User or client devices can include any of a number of general purpose personal computers, such as desktop or laptop computers running a standard operating system, as well as cellular, wireless, and handheld devices running mobile software and capable of supporting a number of networking and messaging protocols. Such a system also can include a number of workstations running any of a variety of commercially-available operating systems and other known applications for purposes such as development and database management. These devices also can include other electronic devices, such as dummy terminals, thin-clients, gaming systems, and/or other devices capable of communicating via a network.

Most embodiments utilize at least one network that would be familiar to those skilled in the art for supporting communications using any of a variety of commercially-available protocols, such as Transmission Control Protocol/Internet Protocol (TCP/IP), File Transfer Protocol (FTP), Universal Plug and Play (UPnP), Network File System (NFS), Common Internet File System (CIFS), Extensible Messaging and Presence Protocol (XMPP), AppleTalk, etc. The network(s) can include, for example, a local area network (LAN), a wide-area network (WAN), a virtual private network (VPN), the Internet, an intranet, an extranet, a public switched telephone network (PSTN), an infrared network, a wireless network, and any combination thereof.

In embodiments utilizing a web server, the web server can run any of a variety of server or mid-tier applications, including HTTP servers, File Transfer Protocol (FTP) servers, Common Gateway Interface (CGI) servers, data servers, Java servers, business application servers, etc. The server(s) also may be capable of executing programs or scripts in response requests from user devices, such as by executing one or more Web applications that may be implemented as one or more scripts or programs written in any programming language, such as Java®, C, C# or C++, or any scripting language, such as Perl, Python, PHP, or TCL, as well as combinations thereof. The server(s) may also include database servers, including without limitation those commercially available from Oracle®, Microsoft®, Sybase®, IBM®, etc. The database servers may be relational or non-relational (e.g., “NoSQL”), distributed or non-distributed, etc.

The environment can include a variety of data stores and other memory and storage media as discussed above. These can reside in a variety of locations, such as on a storage medium local to (and/or resident in) one or more of the computers or remote from any or all of the computers across the network. In a particular set of embodiments, the information may reside in a storage-area network (SAN) familiar to those skilled in the art. Similarly, any necessary files for performing the functions attributed to the computers, servers, or other network devices may be stored locally and/or remotely, as appropriate. Where a system includes comput-

erized devices, each such device can include hardware elements that may be electrically coupled via a bus, the elements including, for example, at least one central processing unit (CPU), at least one input device (e.g., a mouse, keyboard, controller, touch screen, or keypad), and/or at least one output device (e.g., a display device, printer, or speaker). Such a system may also include one or more storage devices, such as disk drives, optical storage devices, and solid-state storage devices such as random-access memory (RAM) or read-only memory (ROM), as well as removable media devices, memory cards, flash cards, etc.

Such devices also can include a computer-readable storage media reader, a communications device (e.g., a modem, a network card (wireless or wired), an infrared communication device, etc.), and working memory as described above. The computer-readable storage media reader can be connected with, or configured to receive, a computer-readable storage medium, representing remote, local, fixed, and/or removable storage devices as well as storage media for temporarily and/or more permanently containing, storing, transmitting, and retrieving computer-readable information. The system and various devices also typically will include a number of software applications, modules, services, or other elements located within at least one working memory device, including an operating system and application programs, such as a client application or web browser. It should be appreciated that alternate embodiments may have numerous variations from that described above. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, software (including portable software, such as applets), or both. Further, connection to other computing devices such as network input/output devices may be employed.

Storage media and computer readable media for containing code, or portions of code, can include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information such as computer readable instructions, data structures, program modules, or other data, including RAM, ROM, Electrically Erasable Programmable Read-Only Memory (EEPROM), flash memory or other memory technology, Compact Disc-Read Only Memory (CD-ROM), Digital Versatile Disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by a system device. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

In the preceding description, various embodiments are described. For purposes of explanation, specific configurations and details are set forth in order to provide a thorough understanding of the embodiments. However, it will also be apparent to one skilled in the art that the embodiments may be practiced without the specific details. Furthermore, well-known features may be omitted or simplified in order not to obscure the embodiment being described.

Bracketed text and blocks with dashed borders (e.g., large dashes, small dashes, dot-dash, and dots) are used herein to illustrate optional operations that add additional features to some embodiments. However, such notation should not be taken to mean that these are the only options or optional operations, and/or that blocks with solid borders are not optional in certain embodiments.

Reference numerals with suffix letters may be used to indicate that there can be one or multiple instances of the referenced entity in various embodiments, and when there are multiple instances, each does not need to be identical but may instead share some general traits or act in common ways. Further, the particular suffixes used are not meant to imply that a particular amount of the entity exists unless specifically indicated to the contrary. Thus, two entities using the same or different suffix letters may or may not have the same number of instances in various embodiments.

References to “one embodiment,” “an embodiment,” “an example embodiment,” etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

Moreover, in the various embodiments described above, unless specifically noted otherwise, disjunctive language such as the phrase “at least one of A, B, or C” is intended to be understood to mean either A, B, or C, or any combination thereof (e.g., A, B, and/or C). As such, disjunctive language is not intended to, nor should it be understood to, imply that a given embodiment requires at least one of A, at least one of B, or at least one of C to each be present.

The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the disclosure as set forth in the claims.

What is claimed is:

1. A computer-implemented method comprising:

receiving a reference to a data store and credentials to access the data store;

receiving text data from the data store, the text data including a plurality of objects;

receiving an augmented manifest for the text data, each object from the plurality of objects corresponding to a separate line in the augmented manifest, by:

receiving a plurality of files of a data type from the data store, the plurality of files associated with the plurality of objects;

extracting metadata from the plurality of files; and generating the augmented manifest in the data store based on the metadata;

executing a labeling job on a subset of the plurality of objects, the subset of the plurality of objects identified by running a query on the augmented manifest, the labeling job identifying a plurality of labels associated with the subset of the plurality of objects;

adding the plurality of labels to the augmented manifest to generate a new augmented manifest, each label added in-line with its corresponding object, the augmented manifest to be used as input to a second labeling job; storing an output dataset and the new augmented manifest in an output data store; and

incrementally training an active learning service model using the output dataset, wherein the active learning service includes a validation dataset, and sends feedback on the validation dataset to a dataset service to remove particular objects from the validation dataset.

2. The computer-implemented method of claim 1, wherein the new augmented manifest includes the plurality of labels associated with the labeling job and a reference to the plurality of objects in the output data store.

3. The computer-implemented method of claim 1, wherein each object from the plurality of objects is defined in the augmented manifest without using a preamble or post-script to define boundaries of the object in the augmented manifest.

4. A computer-implemented method comprising:

receiving a plurality of objects using a reference to a data store;

receiving an augmented manifest for the plurality of objects, each object from the plurality of objects corresponding to a separate line in the augmented manifest, by:

receiving a plurality of files of a data type in the data store, the plurality of files associated with the plurality of objects;

extracting metadata from the plurality of files; and

generating the augmented manifest in the data store based on the metadata;

receiving a subset of the plurality of objects to be labeled;

executing a labeling job on the subset of the plurality of objects, the labeling job identifying a plurality of labels associated with the subset of the plurality of objects and

generating a labeled subset of the plurality of objects; updating the augmented manifest based on the plurality of labels identified in the labeling job, each label added to a line of the augmented manifest that is associated with an object from the subset of the plurality of objects corresponding to the label, the augmented manifest to be used as input to a second labeling job; and

incrementally training an active learning service model using the labeled subset of the plurality of objects, wherein the active learning service includes a validation dataset, and sends feedback on the validation dataset to a dataset service to remove particular objects from the validation dataset.

5. The computer-implemented method of claim 4, wherein receiving a reference to a data store further comprises:

receiving credentials to read the plurality of objects and the augmented manifest, wherein the obtaining of the augmented manifest is based on use of the credentials.

6. The computer-implemented method of claim 4, wherein receiving a subset of the plurality of objects to be labeled further comprises:

executing a query on the augmented manifest to identify the subset, wherein the query includes a filtering condition executed on the augmented manifest to filter the plurality of objects.

7. The computer-implemented method of claim 4, wherein receiving a subset of the plurality of objects to be labeled further comprises:

randomly sampling the plurality of objects using a large dataset sampling technique.

8. The computer-implemented method of claim 7, further comprising:

combining the plurality of objects and the plurality of labels into a new output dataset; and

storing the new output dataset and a new augmented manifest in an output data store.

9. The computer-implemented method of claim 4, further comprising indexing each line of the augmented manifest.

23

10. The computer-implemented method of claim 9, wherein updating the augmented manifest based on the plurality of labels identified in the labeling job, further comprises:

identifying a line of the augmented manifest corresponding to a first label from the plurality of labels; and
writing the first label to the line of the augmented manifest corresponding to the first label using a byte offset indexed for the line of the augmented manifest corresponding to the first label.

11. The computer-implemented method of claim 4, wherein each object from the plurality of objects is defined in the augmented manifest without using a preamble or post-script to define boundaries of the object in the augmented manifest.

12. The computer-implemented method of claim 4, wherein the plurality of objects includes a reference to at least one of a video, an image, a multi-dimensional vector, text data, or a human-generated label.

13. A system comprising:

a data labeling service implemented by a second one or more electronic devices, the data labeling service including instructions that upon execution cause the data labeling service to:

receive a plurality of objects using a reference to a data store;

receive an augmented manifest for the plurality of objects, each object from the plurality of objects corresponding to a separate line in the augmented manifest, by:

receiving a plurality of files of a data type from the data store, the plurality of files associated with the plurality of objects;

extracting metadata from the plurality of files; and
generating the augmented manifest in the data store based on the metadata;

receive a subset of the plurality of objects to be labeled; execute a labeling job on the subset of the plurality of objects, the labeling job identifying a plurality of labels associated with the subset of the plurality of objects and generating a labeled subset of the plurality of objects;

update the augmented manifest based on the plurality of labels identified in the labeling job, each label added to a line of the augmented manifest that is associated with an object from the subset of the plurality of objects corresponding to the label, the augmented manifest to be used as input to a second labeling job; and

24

incrementally train an active learning service model using the labeled subset of the plurality of objects, wherein the active learning service includes a validation dataset, and sends feedback on the validation dataset to a dataset service to remove particular objects from the validation dataset.

14. The system of claim 13, wherein to wherein a subset of the plurality of objects to be labeled, the instructions, when executed, further cause the data labeling service to:

execute a query on the augmented manifest to identify the subset, wherein the query includes a filtering condition executed on the augmented manifest to filter the plurality of objects.

15. The system of claim 13, wherein to receive a subset of the plurality of objects to be labeled, the instructions, when executed, further cause the data labeling service to:

randomly sample the plurality of objects using a large dataset sampling technique.

16. The system of claim 15, wherein the instructions, when executed, further cause the data labeling service to:

combine the plurality of objects and the plurality of labels into a new output dataset; and

store the new output dataset and a new augmented manifest in an output data store.

17. The system of claim 13, wherein the instructions, when executed, further cause the data labeling service to: index each line of the augmented manifest.

18. The system of claim 17, wherein to update the augmented manifest based on the plurality of labels identified in the labeling job, the instructions, when executed, further cause the data labeling service to:

identify a line of the augmented manifest corresponding to a first label from the plurality of labels; and

write the first label to the line of the augmented manifest corresponding to the first label using a byte offset indexed for the line of the augmented manifest corresponding to the first label.

19. The system of claim 13, wherein each object from the plurality of objects is defined in the augmented manifest without using a preamble or post-script to define boundaries of the object in the augmented manifest.

20. The system of claim 13, wherein the plurality of objects includes a reference to at least one of a video, an image, a multi-dimensional vector, text data, or a human-generated label.

* * * * *